PURPOSIVE SYSTEMS

THEORY AND APPLICATION

(FINAL REPORT)

TECHNICAL DOCUMENTARY REPORT NO. ESD-TDR-64-

JULY 1964

H. Edward Massengill, Jr.

DECISION SCIENCES LABORATORY
DEPUTY FOR ENGINEERING AND TECHNOLOGY
ELECTRONIC SYSTEMS DIVISION
AIR FORCE SYSTEMS COMMAND
UNITED STATES AIR FORCE
L. G. Hanscom Field, Bedford, Massachusetts

Project 4690
Task 469003

AD609377

ERRATA

In reference to ESD-TDR-64-531, please note the following:

1. That conflicting notices appear on the inside front cover. Attached
    hereto is a list of the correct notices.

2. Pages 83 and 84 were omitted from some copies of the report and are
    attached hereto.

The following notices apply to this report, all others should be disregarded:

PURPOSIVE SYSTEMS

THEORY AND APPLICATION

(FINAL REPORT)

TECHNICAL DOCUMENTARY REPORT NO. ESD-TDR-64-

JULY 1964

H. Edward Massengill, Jr.

DECISION SCIENCES LABORATORY
DEPUTY FOR ENGINEERING AND TECHNOLOGY
ELECTRONIC SYSTEMS DIVISION
AIR FORCE SYSTEMS COMMAND
UNITED STATES AIR FORCE
L. G. Hanscom Field, Bedford, Massachusetts

Project 4690
Task 469003

This contract work was begun April 1, 1963 and continued through
June 30, 1964. The first six months of the work was done at the Institute
for Research in State College, Pa. while the last nine months was done at
a branch of the Institute for Research in Medford, Massachusetts.

Dr. Anne W. Story was the contract monitor. The contract work was under
the supervision of Dr. Emir H. Shuford, Jr., now at the Decision Sciences
Laboratory in Bedford, Massachusetts. Dr. Shuford served as principal
investigator from April through August of 1963. During this period,
Dr. Masanao Toda, now a member of the Department of Psychology at the
University of Hokkaido in Sapporo, Japan, worked closely with Dr. Shuford
in developing the approach explicated in this paper. Mr. Jun-ichi Nakahara,
now also a member of the Department of Psychology at the University of Hokkaido,
worked under the contract with Drs. Shuford and Toda in developing the
fungus-eater approach. Mr. Edward Massengill, now at the Decision Sciences
Laboratory, worked under the contract from its inception and became principal
investigator in September of 1963 upon the resignation of Dr. Shuford.
Mr. Samuel Vaughn, now with HRB-Singer, Inc. at State College, Pa., served
as administrative assistant on the contract from its inception through
February of 1964.

This final report, summarizing the contract research, developed out of
a study of the contract reports by the author and out of discussions of the
contract work between the author and Dr. Shuford. Thus, the basic ideas
found in the paper are drawn from the contract reports and this paper is
merely an attempt to put them into the context of the approach and theory

which have guided and developed out of the contract research. The particular ideas included and the emphasis given to these ideas are largely the responsibility of the author. The paper, as a whole, represents an interpretation by the author of ideas in the contract reports. Of the other three professionals, besides the author, who worked under the contract, only Dr. Shuford has had a chance to read the paper and make suggestions.

PURPOSIVE SYSTEMS

THEORY AND APPLICATION

(FINAL REPORT)

ABSTRACT

The purpose of this paper is to summarize the approach and theory on which the research performed under ESD contract AF 19(628)-2968 is based. Basically, the approach is the use of decision theory, with the assumption that people behave optimally given their formulations and constraints, to study the significant tasks that people perform. The ultimate goal of the approach is to map human behavior onto logic and mathematics.

The emergence of the approach is given along with four basic requirements that we make of any theory to be used in understanding the behavior of individuals. Our approach is contrasted with more traditional approaches. The procedure of our approach, task analysis, is explained and is illustrated by examples from the contract research. The place of applications in our approach is dealt with extensively.

The paper includes a guide to the more important ideas dealt with in the contract research with references to the relevant contract publications. Abstracts of these publications, seven completed and seven in preparation, are also included.

PUBLICATION REVIEW AND APPROVAL

This Technical Documentary Report has been reviewed and is approved.

DONALD W. CONNOLLY
Acting Chief, Decision Techniques Division
Decision Sciences Laboratory

ROY MORGAN
Colonel, USAF
Director, Decision Sciences Laboratory

TABLE OF CONTENTS

PURPOSIVE SYSTEMS

THEORY AND APPLICATION

(FINAL REPORT)

I. INTRODUCTION

1. Basic research and applications

The study of human behavior involves both basic research, the attempt
to develop a science of behavior, and applied research, the attempt to
use what is presently known about human behavior to help man deal more
successfully with his problems. Thus, basic research leads to the discovery
of principles of behavior, while applied research makes use of these
principles. Basic research and applied research are complementary. The
results of the search for principles of behavior can lead to suggestions
of various ways of improving present practice. And, of course, the very
existence of problems often leads to attempts to find principles to explain
them.

This interaction between the basic and applied aspects of a science
is certainly not limited to psychology. Physicists and engineers have not
waited for a complete understanding of the physical aspects of the universe
in order to apply what they already know. Physicians have not been content
to wait until all of the theoretical issues concerning disease and health
have been solved before applying the principles that are known at any time,
however tenuous they may be. The increase in conveniences and in life
expectancy and general health, where these applications have been made,
points up how valuable applications can be even when they come before all
of the principles of an area are known.

This attempt to apply what is known at any stage seems to be true of science in general. Man wants to understand how the universe and the things in it operate. But since he must live in the universe, this attempt at understanding also includes the idea of learning to better control his surroundings in an effort to improve the conditions of living. But even if man were only interested in understanding the universe, it seems clear that to do a thorough and efficient job, he would need to evaluate, by application, the principles he believed to be true and to build on those that seem to be effective. For example, progress in understanding and discovering principles of flight has no doubt been accelerated as a result of the application of the principles known at any point in developing means of flight.

It should be clear that both basic and applied research have an important place in science in general and psychology in particular. Man's life has been greatly affected by application of what is known about human behavior. Certainly the work done concerning mental illness has led to an improvement in the lot of those who are placed in this category. Again, research concerning exceptional children has been effective in helping children, who in previous times would have been regarded as hopeless, to gain skills that will enable them to lead more independent lives. And much of the progress that has been made in applications has been the result of prior theorizing and experimentation.

But the mere fact that principles have been and are being discovered and applied in an area, does not guarantee that the approach being used is leading to the discovery of the most important principles involved or that it will ever do so. There are a very large number of possible relations existing at any one time in a particular scientific area and it

is impossible to examine them all. Thus, we as scientists must continually evaluate the classes of principles we are trying to discover in terms of their importance, if it is our purpose to discover the significant relations in our area. We must evaluate our underlying theory and approach in terms of how well they are helping us to deal with the important relations. This includes a comparison of our approach with other approaches in an effort to determine which will be more effective, in terms of leading to the discovery of important relations.

2. The emergence of our approach

The work, described in the abstracts in chapters V and VI, is the result of a definite view concerning the theory and approach that should be used in developing a science of behavior and the approach that should be taken in applications. This view has developed as a result of two processes: the development of decision theory and a growing dissatisfaction with current theories and approaches in psychology. Of course, there has been interaction between these two processes. The development of decision theory, its promise in leading to a science of behavior and its contributions to important applications, has contributed to our dissatisfaction with traditional theory and approaches. And this dissatisfaction has, in turn, encouraged greater effort in developing decision theory.

Our goal in developing a science of behavior is to understand the behavior of individuals. It seems to us that this can best be accomplished by studying the systems involving individuals and the tasks they perform. It is our assumption that the performance of a task by an individual is guided by his purpose in performing the task. Hence we regard these systems as purposive systems. Having gone this far, we need a theory which will enable us to effectively and efficiently study these purposive

systems. Because of our particular scientific backgrounds, there are certain characteristics that we desire in a theory that is to be used in studying purposive systems.

1. It should be a comprehensive theory.

2. Its language should be a purposive language.

3. It should be internally consistent.

4. We should be able to map behavior onto the theory.

3. Discussion of our requirements for a theory

As we will emphasize throughout this paper, we want to begin with a comprehensive theory. We are not interested in studying isolated aspects of behavior with a view of someday trying to fit them together. Rather, we want to start in the beginning with a theory that will be able to handle, in principle, all aspects of behavior. We have assumed that the systems we desire to study are purposive and thus we want our theory to be stated in the language of purpose (Toda and Shuford, 1964b). We want the theory to be above reproach in its relation to logic, therefore we insist that it be internally consistent. And finally, we want to be able to map behavior onto the theory once it has been developed. To do this, it will be necessary for the statements of the theory to be such that a one-to-one mapping with behavior is possible. And to do the mapping, we will have to make some assumption(s) that will connect behavior with the theory. We would like for this assumption or these assumptions to sound as reasonable as possible at the time we begin to use the theory, so that we will have some confidence that development of the theory will yield the desired results.

We have, in decision theory, a theory that fulfills the above requirements. We should not leave the impression here that we specifically

set up these criteria before we chose to use decision theory in our work. It is closer to the truth to say that the use of decision theory has pointed up the importance of these criteria. But let us look more closely at how decision theory meets these requirements.

Comprehensiveness. Decision theory is a comprehensive theory in that, if applicable to behavior, it can explain all aspects of behavior. This can be compared to a more restricted theory that might purport to explain only the learning of nonsense syllables. And as a comprehensive theory, decision theory offers an economy of description for behavior as a whole. There are many ways in which behavior could be described. There are many levels on which the description might take place. Decision theory explains behavior on a fairly macroscopic level as opposed, for example, to a theory which might deal with behavior in terms of atoms. The very fact that decision theory operates at a macroscopic level means that the total number of possible relations with which we will be concerned is greatly reduced.

Purposive language. The language of decision theory is the formal statement of the informal language that human beings have developed for talking about behavior in the type of tasks that we regard as important in developing a science of behavior.[1] For example, we talk about "why" we do certain things, the "chance" of "a certain thing happening," how much something is "worth" to us, the "result" of taking a particular "course of action" when a particular state of the world is the case, the "cost of finding out more" about a situation, how much we can expect to "profit" by following one course of action as opposed to another, the "best course of action" to follow. Persons familiar with decision theory

---

[1]See Toda and Shuford (1964b).

will immediately recognize the language equivalents of the above ideas in decision theory: purpose, probability, state of nature, utility, outcome, act, cost of information, expected gain, optimal act. Human beings have observed themselves and others performing tasks and have used words to describe what they perceived. Decision theorists have merely set up a formal language that expresses these aspects of purposive tasks.

In setting up this formal language, an economy in expression has been brought about. In our everyday language, we use several different terms to refer to the same thing or one term to refer to several different things. For instance, we talk about probability, chance, likelihood. Decision theory uses just one term to stand for this concept and that is probability. Thus, the number of words that need to be used in discussing purposive systems are sharply reduced by using the language of decision theory since each concept has but one term associated with it and each term has but one concept associated with it.

Internal consistency. Decision theory is internally consistent. The internal consistency of decision theory is assured because decision theory is essentially the incorporation of logic and mathematics to find maxima or minima. In this sense, decision theory is simply applied mathematics. As we will point out later, decision theory specifies the maximum expected utility or minimum expected loss given the formulation and constraints of the problem in question. Since, in the language of decision theory, the act with the maximum expected utility or minimum expected loss is defined as the optimal act, decision theory specifies optimal behavior given the formulation of the problem and the constraints involved. If one accepts logic and mathematics, he will agree, if he is consistent, that decision theory does specify optimal behavior in the above sense. Thus, it is evident that decision theory is internally consistent.

The mapping of behavior. Finally, decision theory yields statements that have a one-to-one correspondence with behavior. And the mapping of behavior onto the theory can be done with one assumption. This assumption seems to be a reasonable one on the basis of what we know about behavior. The assumption is that people behave optimally given their formulations and constraints. In other words, a person behaves in any situation so as to maximize subjectively expected utility, given his formulation of the situation and the constraints involved. We will comment more on this assumption in section 2 of chapter II.

At this point, we should make clear that while the assumption sounds reasonable to us, we realize that there is a chance that it is not the case, or at least that it is not exactly the case. We are not going to say definitely that people do behave this way. We do not have the evidence necessary to make this statement. But because we have good reason to believe that they do and because making this assumption will allow us to map behavior onto logic and mathematics, by allowing us to map it onto our theory, we make it an assumption in our approach. Since it seems to be a reasonable assumption, we have confidence that our approach to developing a science of behavior will be effective.

Our assumption seems to be a relatively weak one to have to make in order to get a mapping which promises such powerful results. We could have made a much stronger assumption, e.g., that people behave optimally given some decision theorist's formulations of their tasks. Using this assumption, we would probably arrive more quickly at a point where the theory could be tested. But this assumption does not seem very reasonable to us and we believe that we could quickly cite examples in which it is contradicted. On the other hand, by using the assumption that we are

willing to make, it will be a long time before anyone is able to find a contradictory example.[1] This matter is discussed further in section 4 of chapter II.

Thus, we have an assumption, relatively weak and relatively reasonable, in terms of what we know about human behavior, which will enable us to map behavior onto logic and mathematics. This mapping would not be possible, however, even given the assumption that we make, if the statements of the theory were not such that a one-to-one mapping could be made. But fortunately, decision theory yields such statements. There is a kind of uniqueness in the solutions of decision theory that makes a one-to-one mapping possible. This distinguishes the results of decision theory from the results of the satisficing theory. The criterion in this latter theory is to choose the act which has an expected utility greater than some given value. It is easy to see that we could not get the desired mapping in this case because there is no one best act as there can be in decision theory, i.e., in the satisficing theory, acts with different expected utilities can meet the criterion.

Thus, decision theory seems to be the ideal theory for use in the study of purposive systems when the goal is to understand the behavior of individuals. In the next section, we will contrast our approach with other approaches. In doing so, we will expand on the ideas that we have introduced in this chapter.

---

[1]For an illustration of how difficult one is to find, see the section of Toda and Shuford's paper (1963) concerning the urn gamble.

II. OUR APPROACH CONTRASTED WITH OTHER APPROACHES

This chapter deals with eight aspects which show how our approach compares with other approaches. These aspects are neither mutually exclusive nor exhaustive. They are:

1. The type of theory: comprehensive versus miniature theories.

2. The content of the theory.

3. The subject matter of psychology.

4. The place of empirical research.

5. The type of tasks to be studied.

6. The task performers of interest.

7. The individual versus the "average" individual.

8. The quality of performance.

1. The type of theory

In the earlier days of psychology, comprehensive theories played an important role in research. The best known of these theories are usually grouped under the headings of stimulus-response theories and Gestalt theories. But of late, there has been a tendency in psychology to forego the comprehensive theory in favor of miniature theories, i.e., theories that cover very restricted areas of behavior.

In our opinion, there is more to be gained by beginning with a comprehensive theory and developing it to take into account various aspects of behavior than by beginning with isolated theories and attempting to put them together into a unified whole. One reason is that we are not convinced that one can necessarily build higher level relations from more

elementary relations. In other words, we are not sure that miniature theories can be combined to yield a comprehensive theory. Another reason is that valuable time can be wasted by concentrating on irrelevant aspects of behavior, i.e., developing irrelevant miniature theories.[1]

2. The content of the theory

Though we favor comprehensive theories, we find that we cannot accept the stimulus-response approach because of its seeming inability to handle meaningfully the high level tasks in which we are interested. And while we feel more at home with the Gestaltists, when we hear such words as structure and wholes, we feel that they have not sufficiently developed their ideas. Thus we have turned to another comprehensive theory, decision theory. As we have said, decision theory is of use to us because it is a comprehensive theory, because its language enables us to talk rigorously about purposive systems, because it is internally consistent, and because, with the addition of the rationality assumption, it allows us to infer the formulations and constraints of individuals.

In the language of decision theory, our assumption is that individuals maximize subjectively expected utility, given their formulations and constraints. Note that this is not an assumption of decision theory, as such, but merely a use of its language. Decision theory merely yields the maximum expected utility or minimum expected loss given the probabilities, utilities, purpose, and constraints involved in a situation. And this is just the result of applying logic and mathematics to find this maximum or minimum expectation. Decision theory itself does not comment on whether

---

[1]This question of comprehensive versus miniature theories is examined further in section 4 of this chapter.

a person should or does behave in a certain way. It merely specifies optimal behavior for the given situation.

But our rationality assumption plus decision theory, enables us, at least in theory, to look at observed or possible behavior and explain this behavior by specifying the formulation and constraints of the person involved. In actual practice, it will depend on the extent to which decision theory has been developed and, of course, on how well our assumption holds in the situation in question. But if our assumption does hold, then development of decision theory will lead to an understanding of the behavior of individuals in the sense that formulations and constraints can be inferred on the basis of observed behavior.

At this point, we should comment further on what we mean by saying that people behave optimally given their formulations and constraints. This implies that different people can perform the same task with different resulting behavior and yet the behavior of each be optimal. There are two important points to be considered here. One is that two people, or one person at different times, may perceive highly similar situations differently. It may be because of differences in perceptual ability, because of a choice to ignore certain factors, etc. The other point is that two people who perceive a situation in the same way may make different use of the information that is present because of a difference between the decision principles that they have acquired, because of constraints in their capacities, etc. But we assume that when formulations and constraints are taken into consideration, the resulting behavior is optimal. Thus, when decision theory is more fully developed, we should be able to infer, on the basis of observed behavior, the formulation and constraints involved in a given purposive system.

3. The subject matter of psychology

In our view, psychology is concerned with individuals performing tasks. The performer and his task can be regarded as a system. Thus the subject matter of psychology is these systems. Most psychologists would probably agree with us up to this point. But we go one step further and describe these systems as purposive systems. Thus, in our view, the subjective matter of psychology is purposive systems.

This, of course, clashes with the ideas of those psychologists who would regard the systems as basically mechanistic. For instance, an S-R psychologist might say that the use of purpose is superfluous; that the systems can be explained in terms of stimuli evoking responses. But it is our view that in high level tasks, the performer chooses the response(s) that will result from a given set of stimuli. We assume that the person in the system is reacting to stimuli by attempting to organize the incoming information. This organizing involves judging the uncertainty of the possible states of the world, placing values on possible outcomes, etc. And further, this organizing is being done in order that the course of action with the highest subjectively expected utility may be found. Thus, in our view, the person in the system is very active in determining the response(s) that will be made to the stimuli.

4. The place of empirical research

We have said that decision theory specifies optimal behavior, i.e., the act with the maximum expected utility or minimum expected loss, given a well-specified task, including the constraints involved. This is decision theory in its normative sense. And decision theory is truly normative in this sense. In other words, the behavior it specifies is optimal and anyone: animal, man, or machine, who has the same formulation and the same

constraints, should behave in the manner prescribed by decision theory if he wants to behave optimally. This just means that given the logic we accept, the behavior specified by decision theory is optimal. In other words, it is the consequence of using logic and mathematics to find maxima or minima. Thus, given a formulation, a purpose, and the constraints involved, decision theory does specify the optimal behavior.

We should hasten to add that we cannot chose a task, formulate it, find the optimal behavior, and then say that everyone should behave this way in order to be optimal. This is not a consequence of logic. As we have said, different people may formulate the same problem in different ways or may consider different aspects of the same problem. Thus what is optimal for one person may not be optimal for another, even in the same situation. So the problem must be fully specified before the should of decision theory can really have meaning.

So first, decision theory is normative in the above sense, i.e., for a well-specified problem, the behavior prescribed by decision theory is optimal; logic and mathematics work. And second, we assume that when the task of a person is fully specified, i.e., the formulation, purpose, and constraints known, his behavior will be found to be optimal. In other words, we assume that it conforms with decision theory and thus with logic and mathematics. Given these two points, it is possible to understand the behavior of individuals by studying the formulations of tasks that would result given certain constraints, or given that certain constraints were removed, and by noting the behavior that results. When this has been done for a very large number of significant tasks, it should be possible to look at human behavior and infer the formulations and constraints that produced

it. Thus, we can develop our theory by logical deduction rather than by empirical observation.

There seem to be three possible uses of empirical observation in developing a science. First, empirical observation may help to suggest a given theoretical approach. Once a theory has been adopted for use in understanding an area, it may be necessary to continue to use empirical observation in order to extend it. This implies a certain weakness about the theory in that one would like to be able to develop the theory, once the assumptions have been made, by logical deduction. Once a theory has been developed, there is a need to test its effectiveness in explaining the area it purports to explain. This state, of course, involves empirical observation.

We have, in our work, already dealt with empirical observation in the first sense. We have observed human behavior and studied the findings of others and have, on this basis, formulated a seemingly fruitful approach to further study. Now our job is to develop the theory. Since our theory is so structured as to allow us to develop it by logical deduction, we have no need of additional empirical observation in developing it. We need only study tasks. There will come a day when we will wish to test the theory, but this day seems at present to be far off, i.e., our theory will not be testable until it has been developed to a much greater extent.

Of course, our emphasis on the development of decision theory by the analysis of tasks does not mean that we will no longer do any empirical work. In fact, as we analyze more tasks, we may find that we need to utilize empirical research in ways that we can't imagine now. But for the present, our empirical work will be in the area of applications. And decision theory has very important applications, e.g., helping people to

learn to formulate their important decision problems so as to increase expected payoff; helping them to use their own utilities, probabilities, and formulations to get exact solutions for their problems; and furnishing them with approaches for handling problems that have not yet been analyzed. Thus while empirical work is not necessary at this time in developing our theory, it has an important role in the area of applications.

Of course, there will be many who will not agree that we should proceed without further empirical observation. This probably means that they do not agree with our rationality assumption and/or with our assumption that decision theory can be developed so as to enable us to understand behavior. Many of those that disagree probably prefer to keep close to the empirical and to deal with theories that are restrictive enough so as to be amenable to testing immediately.

While many psychologists will distrust an approach which is not immediately amenable to testing, we feel that there is ample reason to distrust the prevalent approach to developing miniature descriptive theories. We do not deny that these theories may actually describe, in the relevant population, the aspect of behavior they are meant to describe. But because of the important part that culture and learning play in determining behavior, it seems clear that the empirical study of behavior is apt to involve arbitrariness and instability.

The fact that the theories deal only with isolated aspects of behavior, means that in order to understand all aspects of behavior, many such theories will be needed. We question the assumption that a large number of these theories, taken together, will lead to an understanding of the important aspects of behavior. We have already seen the difficulty of combining elementary relations in an effort to get more complex ones when

the rules of combination are not known. In the first place, the relations
may change when combined with other relations. In the second place, the
instability and arbitrariness, involved in the systems on which the
miniature theories are based, lead us to believe that any attempt to combine
these theories will result in inconsistencies that will be hard to remove.

Thus in the light of our goal of developing a comprehensive theory of
behavior, we believe that an approach using a theory that is based on the
prior study of human behavior and one which can be developed by logical
deduction has as much to commend it as does an approach based on the continual
observation of human behavior and the use of miniature descriptive theories.

5. The type of tasks to be studied

Traditionally, psychologists have attempted to simplify the systems
they study. This has resulted in the division of psychology into various
areas such as learning, perception, memory, motivation, etc. Going even
further, they have tried to find tasks within these areas which could be
controlled so that, at most, only a few things would vary. They have
introduced the animal into psychological research in an attempt to get task
performers that are easier to control than human beings. The rationale for
studying restricted tasks and simpler organisms is that this will enable
the psychologist to find very elementary relations. And it is believed
that, once enough of these elementary relations have been found, a science
of behavior which encompasses all of them will emerge.

We have grave doubts about many of the tasks currently being studied
in psychology. First, we question the simplicity of many of the tasks
commonly regarded as simple. Second, we question whether some of the tasks
currently being studied are in fact decompositions of any significant tasks
that people perform. And third, we question the relevance of studying these

tasks, in terms of understanding significant tasks, even if they are the results of valid decompositions of higher level tasks.

The complexity of apparently simple tasks. Psychologists try to choose tasks and subjects in such a way as to control most of the aspects of the resulting systems. These resulting systems seem to be simple systems. The fact is that many times they are not well controlled at all and are far from simple.

There are several reasons that an experimenter may regard a task as simple when it is actually very complex. In some cases, he may not understand the task involved to any great degree and thus may believe that it is a simple task when actually it is very complex. But even if he understands the task involved, he may not understand how different subjects will perceive the task. Thus he may decide to ignore various aspects of the situation because they seem unimportant to him. This, of course, does not guarantee that the subjects will regard them as unimportant. If some or all of them do not, then the experimenter is in trouble; because to adequately interpret performance, he needs to know each subject's formulation or be reasonably sure that each subject has the same formulation.

Since we know very little about a subject's formulation in any task, the only way to proceed in order to get meaningful empirical observations is to try to make sure that all of the subjects adopt the formulation and purpose of the experimenter. The success of this approach depends on the experimenter adequately understanding the task, on his being able to communicate the task to the subject, and on his being able to persuade the subject that it is worth his (the subject's) while to use this formulation and purpose. In our empirical research, we have tried to accomplish this by using well-defined tasks and actually making it worthwhile for the

subject to adopt our formulation and purpose. But, of course, most of our empirical work is concerned with applications rather than basic research. In many experiments in basic research, the subject is given the explanation of the task and asked to accept it, with little or nothing in the structure of the task which would induce him to do so. There are many examples of using college sophomores in experiments involving no specific payoff function that might lead them to cooperate. This can turn what seems to be a simple task into a very complex one.

For example, a free recall task seems, at first glance, to be a fairly simple, well-controlled type of task. Subjects are shown a list of words, one at a time, and then are asked to recall as many of them as possible. Let us suppose that the experimenter only asks the subjects to do this; that he does not pay them in such a way that will encourage them to respond with all of the words they can recall and/or will not encourage them to try very hard to remember the words.[1] In this case, the experimenter is asking the subjects to induce a payoff function onto the possible responses, such that the optimal behavior will be to answer with as many responses as can be remembered.

Suppose a subject decides that he would like to appear to be a good learner in the task. He believes that to guarantee this appearance, he should hold some of the words he remembers on the first few trials in reserve, i.e., not respond with them, so he will be sure of having an adequate number of words on the later trials to give the appearance of

_____

[1]Many experimenters deceive subjects about the structure of the task in which the subject is participating. Experimenters should not be so naive as to believe that subjects are not deceiving them, though maybe inadvertently, when they say they understand a task and will try to do their best in the task even when there is little or no incentive to do so.

learning.[1] In this case, the subject is inducing a payoff function over the possible outcomes on each trial that may have a much different maximum than the payoff function which would cause him to respond with all of the words he could remember. Thus he may maximize subjectively expected utility, i.e., behave optimally, but in terms of a payoff function much different from the one that the experimenter desired.

The experimenter has in this situation, by failing to exercise adequate control, left room for the subjects to assume their own payoff functions. At the very least, he could have paid subjects according to the number of responses they gave on a trial and made the payments high enough to encourage the best performance possible from them. This lack of control, an example of which is leaving many things unspecified, opens the door for subjects to supply their own interpretations. The result can be to make a rather simple task very complicated.

In such situations, the observed behavior is probably meaningful only as a commentary on how well a subject was able to do in a specific task, given his own formulation of the task, whatever that formulation was. But it is often taken as an indication of how well he did in the task as formulated by the experimenter; the implication being that he adopted the experimenter's formulation. And, as we shall see in section 7 of this chapter, varied formulations can play havoc with any attempt to average over subjects.

Our lack of understanding of subjects' formulations of tasks and the difficulty of assuring that all subjects in a task adopt the same formulation and purpose, make it very hard to use a simple task in a psychological experiment. This very fact has lead to the use of animals

---

[1]He could complicate the task even more by introducing a measure of uncertainty over the number of words he should hold in reserve.

in an attempt to bypass the difficulties mentioned above. Of course, in
the final analysis, we regard the study of most simple tasks as irrelevant
anyway since we believe that the significant tasks of human beings are the
ones that should be studied.

   Attempt to decompose a significant task into subtasks. There are many
possible ways to analyze a system, involving a human being performing some
task, i.e., to decompose it. It is doubtful that the decomposition of high
level tasks, in terms of the traditional areas of psychology, followed by
the study of various subtasks within an area, will yield the results expected
by traditional psychologists. In the first place, this type of decomposition
by traditional psychologists is not usually applied to some high level task
in particular but to high level tasks in general. For example, memory is
involved in most high level tasks. But memory may play varied roles in
varied tasks and the isolation of supposedly simple tasks involving memory
and the study of these tasks is no guarantee that any relations found will
hold in any important higher level task. Thus, it is sometimes difficult
to picture the high level task that might involve some of the lower level
tasks that have been and are being studied in psychology.

   We should note here that there are genuine relations involved in
these lower level tasks. But we must question their importance in helping
us to understand the behavior of individuals when the tasks are of so little
significance in themselves and when it is doubtful whether they are really
decompositions of significant tasks. We believe that one of the most
important jobs of psychology is deciding which tasks are worthy of study,
in that they will be of most help in developing a science of behavior. In
our view, the answer can only be found by looking for significant, or
crucial, tasks that people perform.

Difficulty of inferring behavior in a complex task from behavior in its subtasks. Even if these So-called simple tasks being studied by many psychologists were the results of valid decompositions of higher level tasks, there is no guarantee that the relations found in systems involving them could be combined to give the relations in the higher level systems composed of them. This is because, in a system involving a higher level task, the interaction between subtasks may invalidate the relations found in systems in which the subtasks were studied separately. In other words, the systems may not be additive.

Let us look at an example which illustrates this point, i.e., that even when there is a valid decomposition of a task into subtasks, the relations found in systems in which the subtasks are studied individually will not necessarily combine to describe behavior in a system involving the original task. Let the original task be a percentage estimation task. Percentages from 0 to 100, i.e., fractions from 0 to 1.0, are generated with probabilities specified by a particular beta distribution. When a percentage, or fraction, p, is generated, it is used in a binomial generating process with a given n to generate r successes and $n-r$ failures. A linear payoff function is used. On a particular trial, a p is generated and used in the binomial process to generate r successes and $n-r$ failures. The subject is asked to estimate the p, knowing the particular prior distribution and the results of the binomial process and using the linear payoff function. The same prior and the same n are used for all of the trials. Notice that both p and r are random variables.

Now suppose that an experimenter decides that he will try to learn about behavior in systems involving the original task by studying behavior in systems involving subtasks of the original task, namely the n, r

combinations. There will be $n+1$ of these combinations. Each combination is produced by one of the possible p's, with a certain probability. Now suppose the experimenter presents the subject with a given n, r combination over a series of trials. The actual p's producing this combination will be generated by the appropriate posterior distribution and thus p will be a random variable. This will be done for each of the n, r combinations and each combination will be handled as a separate task. Each combination will be presented in proportion to the number of times it would theoretically occur in the original task. The experimenter will get responses from each subject for each n, r combination. Then he will attempt to use a given subject's results in the subtasks to explain how that subject would behave in the original task. In other words, he will attempt to combine, for a given subject, the elementary relations found in systems involving the subtasks to describe performance in a system involving the original task.

But the combination of the behavior in these subtasks will not necessarily result in the behavior that would be exhibited in the original task. Remember that, in the original task, r is a random variable, while in the subtasks, it is constant for a given combination. In the original task, the subject is responding in one unified situation. In the subtasks, he is participating in several separate tasks. The experimenter has no way of knowing if composition of the purposive systems involving the subtasks will result in the purposive system involving the original task unless he either determines the formulations involved, so that the validity of such composition can be determined, or runs the subject in both tasks. Of course, the first option is not open to him at present. And if he is only interested in performance in the original task, there is nothing to be gained by running each subject in the subtasks, because clearly, it will not

necessarily give him the information he desires, i.e., he may get irrelevant relations, but it <u>will</u> cost him time and effort.

Seeing the difficulties involved, why do experimenters attempt to use performance in subtasks in order to learn about performance in higher level tasks? The reason is that they believe that, if a performance in a higher level task is too complicated to handle, they can still learn about it by studying relations in subtasks of the higher level task. But aside from the difficulty of isolating these subtasks, it is clear that this approach will not necessarily give the desired results. It is our aim to try to understand behavior in the more significant tasks, by studying the tasks themselves, rather than by trying to isolate and study their subtasks. We will deal with the important tasks that we can specify and we will try to develop our methodology so that we can deal with those that are now out of our reach.

In summary, first, we question whether many of the tasks usually regarded as simple are actually simple. Second, we question whether these tasks are subtasks of significant higher level tasks. And third, even if they are subtasks of higher level tasks, we question the effectiveness of attempting to combine elementary relations found in them in order to understand behavior in the higher level tasks involving them. Our approach is to try to find the crucial, or significant, tasks of human beings and then to determine what optimal behavior results from various formulations given various constraints. The results of this approach should ultimately enable us to look at the behavior of an individual and specify his formulations and constraints.

6. The task performers of interest

Under this heading, we want to discuss the task performers that are of interest to us. In using our rationality assumption to map human behavior onto logic and mathematics, we are basically interested in human behavior. But we would certainly not exclude the behavior of animals. Traditionally, psychologists who have been interested in animal behavior, have been interested mainly because they believed that a study of animal behavior would simplify the matter of finding relations that would be of use in explaining human behavior. We take a different point of view. We believe that a study of tasks may shed light on both human and animal behavior. Since our approach in developing our theory consists solely of task analysis, we will, for the time being, study neither human nor animal behavior. But we hope that our results will ultimately shed light on both types of behavior.

But even if our emphasis were on the study of task performers rather than tasks, we would not study animal behavior in order to understand human behavior. We would look at human performance and try to discover the formulations and constraints that would lead to it. If we did this with animals, it would be because we were primarily interested in animal behavior, not because we thought this would be the most efficient way to learn about human behavior. In other words, if we were going to do empirical studies in an attempt to explain the behavior of human beings, we would study human task performers. If we wanted to explain animal behavior, we would study animal task performers. Even if the study of animal behavior could furnish clues to human behavior, it is not evident that the study of animals in tasks that are important to them would be significantly easier than the study of human beings in tasks that are important to them.

- 24 -

But there is a task performer that we can use in order to help visualize a particular task performer performing the task being studied. In the beginning, the formulations of tasks which we study will be relatively simple compared to the formulations we ultimately want to deal with. It is hard to conceive of animals or human beings performing the well-formulated versions of these tasks because, for the most part, perfect memory and instantaneous calculations are required. And also, the task formulations are so relatively simple that it is hard to imagine human beings who could effectively isolate such small worlds. The answer to having a task performer for which the task formulations would seem plausible is to picture a machine performing the task. This will not only help us in visualizing a concrete task performer performing the task, but will also help us to see how we might complicate the task. We can consider ways of constraining the machine or ways of removing constraints in order to get optimal behavior of various types.

Of course, in applications, we will be interested in human beings, machines, and man-machine combinations. It is not clear, at this point, how interested we will be in applications dealing with animals. Our interest in human beings will mainly concern the improvement of performance. This may involve an increasing cooperation between man and machine in which the machine does what it can do best and the man what he can do best, at any given time. This type of approach is very evident in companies which are beginning to automate. Machines which can do everything have not yet been developed and so machines are used for some jobs and men for others. We believe that it would be well to think very carefully about how man-machine combinations could be used to accomplish some of the jobs that should be done now but which neither man nor machine can do alone.

One example of this is in the area of education. One good tutor teaching one child, can accomplish more, overall, than the best machine setup now in operation. But very few students can be taught in this way because relatively few such tutors are available. Attempts are being made to develop machines which can interact with students in much the same way that a tutor would. But the programming of such machines is a formidable task. It would seem wise to think in terms of what a machine can do best and what a human being can do best, at this time, and attempt to combine the two to get a more efficient and effective teaching system. The machine would make it possible to handle a large number of students, while the human being could be utilized so as to decrease the amount of programming and memory required to interact with students at crucial stages in the process.

A further word is in order concerning the place of machines in applications. Any task that can be well-specified in terms of decision theory, can be performed by a machine. Thus we should be very careful in our thinking not to limit the potential of machines. Our emphasis on task analysis will no doubt open many new areas in which machines will operate more efficiently than men. Since we are concerned with the general improvement of performance, we should be happy to see machines moving into areas in which their use will lead to such improvement.

7. The individual versus the "average" individual

To begin with, let us clarify what is actually included in the term individual. We use this word in the sense of an entity performing a task. It may be a single animal, man, or machine performing a task. It may be more than one animal, more than one man, or more than one machine performing a task. It may be a combination of any of these performing a task. Thus this

entity is not always singular--it can be plural. The point is that the entity, whether one or a group, is viewed as a task performer. Thus when the entity is a group, it is a group performing a task. The members of the group are considered as an entity, though some may be trying to hinder rather than help. Ultimately, we would like to understand as many as possible of the various entities. But in order to make our task of developing a science of behavior fairly manageable, we have chosen to put the emphasis on single human beings, single machines, or man-machine combinations. Throughout this paper, our use of the word individual will primarily apply to these three types of entities, though we don't want to forget the possibility of dealing with others.

And as we come to understand single human beings and machines, we may gain insight into the matter of groups of men and groups of machines performing tasks. As we think in terms of groups, we should be guided by our speculation concerning the combining of elementary relations in order to understand more complex ones. We may suspect that we will not be able to combine the relations found concerning single individuals performing in a task to find those of a group performing in the same task. But the methodology we have developed in studying systems with single individuals may help us in studying systems with groups.

Now that we have shown what the term individual ultimately includes and what we mean to include in it at the present time, let us examine the question of the place of the "average" individual in our work. In our discussion, we will distinguish between two types of averaging. In one, performance is averaged over individuals and over trials to give some grand measure of performance in a task. Statements based on the measure

yielded by this type of averaging usually concern differences between groups from different populations. For example, one group may be trained one way and another group another way. On the basis of the average performances of the two groups in the task, the statement might say how much better an individual drawn at random from one population would be expected to do than an individual drawn at random from the other. Or it might merely say that the difference between the performance of the two would be significant. Obviously this is of little help in explaining the behavior of a given individual, though it might be of some help in predicting behavior. For instance, it is certainly of help to insurance companies.

The other type of averaging has greater potential in helping to explain behavior but we will see that it is actually hard to realize in practice. In this type averaging, the experimenter is trying to discover the process which is generating the observed behavior. He assumes that all of the subjects have about the same formulation of the task and that any difference in performance is due mainly to random error. This method may be effective if the subjects do in fact have the formulations. But as we discussed in section 5 of this chapter, this is very difficult to assure.[1] It is easier to assume in cases where simple tasks are involved. But, as we have noted, many tasks that appear simple are often not simple. And even when simple tasks are obtained and we are fairly sure that the subjects adopt the formulation and purpose of the experimenter, the task may be of little importance in furthering our understanding of behavior.

---

[1] Understanding formulations of subjects seems to be the only way to make success very probable in this type of averaging. And such understanding is not yet at hand.

Thus, we must conclude that averaging over subjects has no place in the development stage of our approach. First, we do not intend to do empirical studies. And, even if we did intend to do them, we have seen above that the first type of averaging would not help us and the second is virtually impossible at this stage. The matter of averaging in applications, however, is another matter.

8. The quality of performance.

In section 4 of this chapter, we have already discussed the matter of optimal performance from the standpoint of basic research. But we believe that it should also be mentioned in terms of applications. In applications, it arises when we seek to comment on the quality of a subject's performance. We will see, in comparing decision theory with traditional approaches, that what can be said by each concerning the quality of a subject's performance is vastly different.

In traditional approaches, the quality of an individual's performance can only be talked about in terms of what perfect performance in the task would be, in terms of the best performance so far as individuals in the task, or in terms of the average performance of a group of subjects in the task. Decision theory, however, allows us to talk about the quality of performance in terms of the maximization of expected utility or the minimization of expected loss. For example, in situations in which the subject is supposed to adopt a given formulation of a problem, we can get some idea of how well he performs by comparing his performance with the strictly optimal performance for that task, where the strictly optimal performance is defined in terms of the experimenter's formulation.[1]

---

[1]See Massengill and Shuford (1964).

We have already assumed that any performance by an individual is optimal given the formulation and constraints involved. So to talk about the quality of performance from this point of view, we must talk in terms of the gain in expected payoff that would result from a different formulation of the task, from the utilization of more of the available information, etc. But we are operating largely in the dark in this type of situation until we better understand the formulations of individuals.

Because of the attempt of traditional psychologists to use simple tasks, it is often possible for a subject to perform perfectly if his capacity is large enough. For example, in free recall experiments, a subject could always respond with all of the words shown to him, if he could store and recall them. When the strictly optimal performance of these tasks is specified without reference to constraints in capacity, it turns out to be perfect performance. But we have already hinted that many tasks, in which being right each time is the strictly optimal behavior, are not very important to human beings. Most of the important tasks human beings perform have some amount of uncertainty involved.

To comment on the quality of performance in higher level tasks, traditional psychologists either have to run subjects in the tasks and get some idea empirically as to the quality of performance possible in the task, or have to try to combine the information they have obtained in systems involving subtasks of the higher level task. By using decision theory, we don't have to depend on perfect performance or on an empirical determination of the quality of performance. Simply by analyzing the task in question, we can determine the optimal performance; optimal in the sense that it is the performance which will lead to the maximum expected payoff or minimum expected loss. Then if individuals are being judged on

the basis of how well they can adopt the formulation, purpose, and constraints of a task, there is a standard, based on logic and mathematics, with which their performance can be compared.

Having this standard, which says that this is the best that can be done, given the information available, the payoff function, and the purpose, can save a great deal of time in deciding which man to choose or whether to choose a man or a machine to perform a task, or whether to try to train a person to perform better in the task. When the individual's performance is compared with the strictly optimal performance, we immediately know how much better the performance could have been and we have a basis for deciding whether or not it would be worthwhile to try to improve it.

Of course, with the ultimate development of a science of behavior, our procedures in applications will be somewhat different. Right now all we can do is to judge the least upper bounds of a person's capacity to adapt to certain well-specified situations. But then, we will be able to infer, from a person's behavior, his formulations and constraints. We will immediately see possible reasons that he did not perform in terms of the strictly optimal specifications and will have a good idea of the steps that need to be taken in order to help him improve his performance.

9. Summary

In summary, we regard our task as the study of crucial purposive systems. We assume that individuals behave optimally given the formulations and constraints involved. Since decision theory offers a comprehensive theory, a rigorous language with which to discuss purposive systems, an internally consistent theory, and allows us to use our

- 31 -

rationality assumption to map behavior onto logic and mathematics, we have chosen it as the basis for our approach.

Because decision theory specifies optimal behavior, given the formulation and constraints in the task and because we assume that people behave optimally given their formulations and constraints, we are able to study purposive systems by analyzing tasks rather than task performers. The most efficient way to reach our goal seems to be through the study of significant tasks rather than through the study of subtasks of these significant tasks. Though our theory may be applied to animals as well as men, we have chosen to begin with men. The practice of averaging over subjects has no place in our approach. And because decision theory specifies optimal performance, we are able to set up meaningful standards of behavior for use in the application of decision theory.

We adopted the idea of systems as purposive systems, the rationality assumption, and decision theory, after due consideration of the current body of knowledge, with which we are familiar, concerning human behavior. This includes the things which we have observed in an uncontrolled fashion throughout our lives. It concludes the results of our empirical research in psychology which seems to suggest that human beings can come close to strictly optimal performance in certain well-specified situations. And it includes a consideration of experimental studies which seem to show that human beings do not behave optimally. Our conclusions concerning studies of this latter type are, of course, that the subjects may not be behaving optimally in terms of the experimenter's formulation but that they may be behaving optimally in terms of their own formulations and constraints.

We see as our next step, in the development of a science of behavior, the analysis of tasks. Our future empirical work will be in the area of

applications, the main purpose of which will be to improve behavior. The next chapter will go more deeply into the matter of task analysis and its place in both basic and applied research.

III. TASK ANALYSIS

We have already stressed the importance of task analysis in our approach. We have also mentioned the type of tasks that we wish to analyze, i.e., the crucial tasks of human beings. In chapter I, we stated the reasons for using decision theory as the basis for our approach. The choice of decision theory, as the underlying theory of our approach, means that our analysis of tasks will be done in terms of decision theory. In other words, decision theory will be used as a tool in the analysis of tasks.

1. Decomposition of tasks

In discussing the analysis of tasks, it is hard to get very far before feeling the need to have a more precise way of discussing what is involved in analysis. This and other considerations led Toda and Shuford (1964a) to investigate the structure of systems. The result was the development of a formal theory of structure. This formal theory provides us with a rigorous way of discussing the analysis and synthesis of systems just as decision theory provides us with a rigorous way of talking about the operation of purposive systems. In the language of structure, analysis is spoken of as decomposition and synthesis as composition.

There is no unique decomposition of a given purposive system, because systems can be decomposed in various ways. For instance, it would be possible for us to use the stimulus-response theory to decompose our systems. Or we might decompose them in terms of the various areas into which psychology has been divided. We might decompose in terms of physiology. Again, we might decompose in terms of atoms. Each of these

decompositions is legitimate. But decision theory is the only decomposition, with which we are familiar, that can fulfill the criteria we have established. Thus, we have chosen decision theory as the basis for our decomposition of purposive systems.

2. The procedure in task analysis

The basic procedure in task analysis is to decompose a task in terms of decision theory. The decomposition tells us whether or not it is a well-specified task and, if not, what else needs to be specified in order to make it one. A task must be well-specified before the optimal behavior associated with it can be found. If the necessary information is not present, we must add it or forget about finding the optimal behavior in the task.[1] Once the task is well-specified, we use decision theory to recompose it; with the result that the optimal behavior is specified. Of course, the optimal behavior was implicit in the task from the beginning but we had to decompose and recompose the task in order to make it explicit. This is analogous to a given area of mathematics where the assumptions implicitly contain all of the implications of the theory but where the consequence of the assumptions must be derived in order to make the implications explicit.

As we said earlier, we plan to proceed by finding the optimal behavior that is associated with various combinations of formulations and constraints in a task rather than to observe behavior and try to specify the formulations and constraints that would make it optimal. To do this, we need to build up a catalogue of formulation-constraint combinations with the resulting optimal behavior of each. Only then will we be ready to attempt an explanation of observed behavior.

---

[1]See Massengill and Shuford (1964) for examples dealing with the use, in applications, of tasks that are not well-formulated.

Thus, our approach will be to choose a task and decompose it. If it turns out not to be well-specified, we will make it well-specified. And we will do so in various ways since there are many possible ways in which a task that is not well-specified can be made well-specified. Once the task is well-specified, we will compose it and note the resulting optimal behavior. But merely finding the optimal behavior for that particular specification will not finish our work with that task. We will add constraints to the task and find the resulting optimal behavior. We will remove constraints and find the optimal behavior that results. We will look at various other formulations that might be made of the task and find the optimal behavior for these. And we will look at behavior in the task that would seem to be inconsistent with logic and try to find the formulation and constraints that would make this type of behavior optimal.

In this way, we will build up the catalogue of formulation-constraint combinations and their resulting optimal behavior patterns. Ultimately, we will obtain enough information so that we can begin to look at an individual's behavior in various tasks and specify the formulations and constraints that are operating. In section 4 of this chapter, we will look at examples of task analysis from the contract publications and, in section 5, we will show how the results of task analysis can be used in applied research. But first, let us examine the matter of the payoff that can be expected from the analysis of tasks.

3. Expected payoff of task analysis

The effectiveness of task analysis in our approach to basic research, i.e., in our approach to developing a science of behavior, depends on our rationality assumption. If it does not hold, then behavior cannot be mapped onto the theory, though we might still profit from the insights

gained concerning behavior. For example, one might not accept Lewin's field theory as being descriptive of behavior but might still gain insight into behavior by studying it.

But we should note that application of the results of task analysis is not dependent on the rationality assumption. In the case of applications, it does not matter whether people do behave optimally given the constraints involved. It does not even matter whether they can do so. Of course, if they can, then the scope of our applications is enlarged and our approach will be different than it would be if they cannot. But we believe that people do behave optimally given their formulations and constraints and can improve their already optimal behavior by adding decision principles, learning to make better use of principles they already possess, etc. If this is true, then the door is open to train people in intuitive decision making and to help them in getting measurements of their probabilities and utilities so that exact solutions to their problems may be obtained.

But even if our rationality assumption does not hold and even if people can't learn to behave optimally, the results of task analysis have application. As we have said, the specification of optimal behavior includes the specifications for achievement of this behavior by machines. This means, at the very least, that we can use the results of task analysis to build machines to perform optimally. And the result will be an improvement in overall performance. Thus, even if our results can only be used by machines, the expected payoff from task analysis is potentially very great. There seems to be little to lose and much to gain by proceeding with task analysis. The end result of task analysis may be a science of behavior. But it definitely will be an improvement in overall performance,

if the results of task analysis are used in applications. Thus, a large payoff can be expected from the analysis of tasks.

4. Examples of task analysis

Now let us look at four examples of task analysis from the contract research forming the basis of this paper. We will look at task analysis in the following contexts:

1. The fungus-eater work involving sequential decision making.

2. Two-act, two-state tasks.

3. Estimation tasks.

4. A two urn gamble problem.

The fungus-eater work. The contract work on the F-E began with the analysis by Toda of four well-defined tasks (1963a), with a constraint on the vision-span, V-span, of the F-E. The length of the V-span was set at 1. Later, this constraint was changed from 1 to 2 and the resulting optimal behavior determined (Nakahara and Toda, 1964b). An informational constraint was applied to the F-E in one of the V-span 1 games, G3. Since, in this game, the F-E can gain information concerning that aspect about which he is uncertain, as he seeks uranium, the resulting optimal behavior takes the form of optimal learning (Nakahara, Shuford, and Toda, 1964a). And finally, the F-E was put into a world involving an accident process which makes the F-E subject to accidental death with a probability known to him. The resulting behavior of the F-E in this larger world was studied (Nakahara, Shuford, and Toda, 1964b).

Two-act, two-state tasks. In the two-act, two-state situation, we have done two types of analysis. One involved learning over decision trials and the other involved the sequential gathering of information

- 38 -

within a trial. In the first case, Shuford (1963) degraded a well-defined situation in four different ways. Because of information available in each situation after the degradation, the resulting optimal behavior was optimal learning over the decision trials. These degradations can be regarded either as resulting from constraints put on the task by a subject or as four different task formulations.[1]

The two-act, two-state well-defined situation was expanded to include a third act which was to take an observation (Massengill, 1964a). In this situation, the observation is generated by a binomial process using one of the two states as a parameter. Since the observations are taken one at a time, with the option to stop at any time, and since there is a limit to the number of observations that may be taken, the resulting situation is a truncated sequential sampling situation. The optimal behavior resulting from the insertion of various values for the prior probabilities, conditional probabilities, and utilities have been obtained (Massengill and Shuford, 1964).

The estimation tasks. The estimation tasks involved a well-defined situation in which each of 101 acts specified the choice of one of 101 states, integer percentages from 0 to 100. The optimal behavior in situations involving differing amounts of additional information was studied. Then certain constraints on information were applied to each situation and the resulting constrained optimal behavior studied. These constraints covered cases in which prior information was ignored, sample information was ignored, both were ignored, and one in which all information,

---

[1]This work was based on the binomial data generating process. Analogous results from the normal and poisson data generating processes are given in (Shuford and Massengill, 1964).

including payoffs, was ignored and responses given randomly. The optimal behavior in each of these situations was compared with the optimal behavior when all the available information was used as well as with the absolute upper limits of performance. This same process was repeated with the purpose to minimize rather than maximize earnings. The matter of optimal fixed sampling was also investigated in this type of task. This work is reported in (Massengill, 1964c; Massengill and Shuford, 1964).

A two urn gamble problem. Notice that the task analysis in the above cases began with a well-defined situation. The optimal behavior in that situation was specified and constraints on information, capacity, or ability were relaxed or imposed. A somewhat different approach was taken by Toda and Shuford (1963). Behavior which was seemingly logically inconsistent was examined in terms of different possible formulations of the task involved and a formulation yielding the observed behavior as optimal behavior was found. It should be noted here that the inconsistencies used were found as a result of empirical observations. But the same type of analysis could be done by asking what if a certain type of behavior were to be observed and proceeding to find the formulation and constraints for which that behavior is optimal.

5. Applications of the results of task analysis

While our purpose is ultimately to be able to understand the behavior of individuals, the results of task analysis can be used in various applications even now. And, as we have said, the application of these results does not depend on our rationality assumption. To illustrate possible applications, we will mention some ways in which the results of task analysis in the contract research, mentioned in section 4 of this chapter, were or could have been applied.

We will be concerned with four types of applications. In one, a person would like to improve his intuitive decision making ability. We could help him in learning to better formulate his problems, in learning to use decision principles with which he is not familiar, and in learning to make better use of the information available in any situation. In a second application, a person would like to get exact solutions for certain decision problems in terms of his probabilities and utilities. In other words, he wants to bypass possible constraints on memory and computational ability and in his ability to use decision principles in combining information. In a third application, an employer would like to hire people who can adopt the company's formulation of a task and perform at a level close to that which would be strictly optimal in terms of the company's formulation. This, of course, may involve both the analysis of performance, in an effort to find people who can perform in this way, and the development of procedures for training people to perform better in this type of situation. In this application, there is the other side of the coin in which potential employees would like to be trained so as to stand a better chance of getting the job. An extension of this application is the design of a decision system in which the task is performed by some man-machine combination and in which each aspect of the task is relegated to the man or the machine according to which can perform it most efficiently at that time. A fourth application concerns the classification of tasks. This involves looking at tasks in terms of some underlying abstract task and comparing tasks in terms of the amount and type of structure they contain.

Intuitive decision making. In the first application, we want to use the results of task analysis to help improve intuitive decision making. In order to do this, it will be helpful to have various formulations of

various tasks in order to show the decision maker how optimal performance changes with changes in formulation and constraints. The methodology that we develop, as we analyze tasks, will be important in helping us to understand the decision maker's formulations and constraints in given tasks and in showing him how these formulations could be improved and the constraints relaxed in order to improve performance. Of course, ultimately, we will be able to observe his behavior in tasks, infer his formulations and constraints and immediately suggest ways of improving his performance. But until we proceed much further, we will have to arrive at formulations by applying the methodology of analysis that has so far been developed as well as by questioning the decision maker and making educated guesses.

One approach that we have taken in seeking to improve intuitive decision making is discussed by Massengill (1964c). We have also attempted to find ways of helping to improve intuitive decision making by studying the various aspects involved in the decision making process. Massengill (1964b) has studied conditional probabilities, in an attempt to better understand the use of empirical and logical conditional probabilities, as well as the processes, both stable and unstable, which produce these probabilities.

Exact solutions. In the second application, we have a person who wants to use his own probabilities and utilities in order to get exact solutions, either to use as such or with which to compare his intuitive solutions. Clearly, in this application, the extent to which decision theory is developed will determine how useful it is in helping such people. And the extent to which the methodology of task analysis is developed will determine how fast exact solutions can be found for specific situations not yet analyzed. And, of course, the further we go in dealing with complicated

- 42 -

problems, the more able we will be to develop heuristic approaches to handle those problems which are still too complicated for our methods.

In making applications of this type, we must be able to handle the formulation for which the person wishes to have an exact solution and maybe help him improve his formulation. But we must also have methods for measuring his probabilities and utilities so that the results of the exact solutions will truly be optimal in terms of his formulation, purpose, probabilities, and utilities. Toda (1963b) describes methods for measuring subjective probabilities in both discrete and continuous situations. Toda and Shuford (1963) examine the concept of utility, the traditional approach to its measurement, and the importance of understanding personal decision contexts in terms of getting valid utility measurements.

Adopting another's formulation. The third application concerns the behavior of a person in a problem in which the formulation and content is specified by someone else. This involves both the judgment performance in such a situation and the training of persons in an attempt to improve performance in the situation, if there is room for improvement. Massengill and Shuford (1964) examine the matter of the analysis of performance in such situations in terms of choosing between men and machines and/or man-machine combinations in both well-formulated situations and situations that are not well-formulated. Two well-formulated situations in which subjects are asked to adopt a particular formulation and content, an estimation situation and a sequential sampling situation, are discussed under the fourth application.

The classification of tasks. A fourth application involves the classification of tasks. There are several ways in which tasks can be classified. We used two of these ways in the contract research: the

classification of tasks in terms of the underlying abstract task and the ordering of tasks in terms of expected payoff. The first, the classification of tasks in terms of the underlying abstract task, can be approached in two ways. One is to begin with a concrete representation and seek to find the underlying abstract task. The other is to begin with an abstract task and give it a concrete representation.

For the task used in the teaching of optimal behavior (Massengill, 1964c), we began with the abstract representation of an estimation task. This task was then interpreted in terms of a concrete business application. It involved the training of prospective employees to estimate the number of top quality pieces in shipments of fruit to a canning company. We could then compare this type of task with actual tasks that are or might be involved in various business applications. In the work on sequential sampling (Massengill, 1964a), a concrete representation of a task was suggested in which an observer was to decide if the blip on a radar scope was caused by a friendly or unfriendly plane. Data could be obtained as the plane came closer. The decision maker could make his decision on the basis of prior information alone or could make it after a given observation of data; up to a certain number of observations. Analysis of this task showed that it was a basically truncated sequential sampling task. The addition of a few constraints gave us an exact concrete representation of this task and thus enabled us to find the strictly optimal performance in the task for various parameter values (Massengill, 1964a; Massengill and Shuford, 1964).

The other type of classification involves the ordering of tasks in terms of expected payoff. Shuford (1963) interpreted the expected payoff of tasks, degraded from a well-defined task, as a measure of the amount of

structure in the tasks. Massengill (1964c) ordered tasks involving

the same conditional information but different prior distributions

according to the worth of the prior information, in terms of expected

payoff.

IV. GUIDE TO IMPORTANT IDEAS DISCUSSED IN CONTRACT REPORT

In this section, we will mention 13 of the more important ideas discussed in the contract publications, with references to the papers in which they are discussed.

1. Constraints

Constraints on information are discussed in the context of two-act, two-state situations by Shuford (1963) and in the context of estimation situations by Massengill (1964c) and Massengill and Shuford (1964). Constraints on capacity, in terms of length of vision-span, are discussed in the fungus-eater, F-E, papers. Nakahara and Toda (1964b) deal with the matter of changing the constraint from a length of 1 to a length of 2. The matter of enlarging a decision context by removing constraints limiting the scope of outcomes is discussed by Toda and Shuford (1963) in the context of a two urn gamble.

2. Constrained optimality

Shuford (1963) differentiates between strict optimality and constrained optimality.

3. Decision contexts

The three decision contexts used to describe the F-E games: the permanent, the external, and the internal contexts, are introduced by Toda (1963a) and discussed in the other F-E papers.

4. Formulations of tasks

Shuford (1963) discusses the formulations of tasks resulting from
the degradation of a well-defined two-act, two-state situation. Various
formulations of the two urn gamble, in which the contents of one urn is
known and the contents of the other is not known, are discussed by Toda
and Shuford (1963). Various formulations of tasks, which result from the
application of constraints to the information available in a percentage
estimation task, are discussed by Massengill (1964c) and by Massengill
and Shuford (1964).

In each of the F-E papers, the formulation of one or more well-defined
tasks is discussed. Formulations for infinite and finite discrete V-span
1 games are discussed by Toda (1963a), with one of the finite games, G4,
being elaborated on by Nakahara and Toda (1964a). The formulation of the
learning F-E in finite game G3 is given by Nakahara, Shuford, and Toda
(1964a). The formulation of the accident-prone F-E in G3 is also discussed
by Nakahara, Shuford, and Toda (1964b). And the formulation of the F-E
with a V-span of 2 is discussed by Nakahara and Toda (1964b).

5. Induced utility

The value associated with an outcome which is a means toward some end
is said to be induced utility rather than utility. This distinction is made
and discussed by Toda and Shuford (1963).

6. Matching property

This concept is discussed by Toda (1963b) and is shown to be of vital
importance in deriving games for the measurement of subjective probability.

7. Measurement of performance

The actual performance of subjects is used as a measure of performance to be compared with the strictly optimal performance in a percentage estimation task (Massengill, 1964c). Nakahara and Toda (1964a) deal extensively with the comparison of actual performance in a finite, discrete F-E game with the optimal performance in that task. Massengill and Shuford (1964) introduce the idea of using the average expected performance of a subject as the measure of performance to be compared with the strictly optimal performance. This concept is applied in the design of a computer program to compare human and optimal performance in the truncated sequential sampling situation (Massengill, 1964b). Massengill and Shuford (1964) also extensively discuss the measurement of performance in applied decision situations.

8. Numerical solutions

Analytic solutions were found for the first three F-E games dealt with by Toda (1963). But subsequent F-E games have required numerical solutions. Numerical solutions are involved for some of the degradations derived by Shuford (1963). Numerical solutions were required for obtaining the expected losses for some of the levels of performance in the percentage estimaton task (Massengill, 1964c). These solutions are also discussed by Massengill and Shuford (1964). A numerical solution is used to find the expected value of the decision tree and the expected optimal sample size in the sequential sampling task (Massengill, 1964b).

9. Optimal learning

Optimal learning in the two-set, two-state situation is discussed by Shuford (1963) for the four degradations he examines. Optimal learning

for finite F-E game G3 is discussed by Nakahara, Shuford, and Toda
(1964a).

10. Strict optimality

Shuford (1963) differentiates between strict optimality and constrained
optimality.

11. Structure

The need for a formal theory of structure is discussed and a formal
theory of structure developed by Toda and Shuford (1964a). This theory
covers the concepts of structure, decomposition and composition, FSS
(formal structure set), closed and open systems, and independence between
systems; among others.

12. Suboptimal procedures

Suboptimal procedures are discussed in two contexts. The first
concerns the selection of suboptimal estimation procedures when the optimal
procedures are not applicable. This is discussed in the second part of
Toda's paper on microstructures (1963a). The second concerns the use of
suboptimal procedures to evaluate machine performance in tasks which are
extremely complex and not well-formulated, so that the optimal performance
cannot be found analytically or numerically, though the basic nature of
the task is known, e.g., dynamic programming, sequential information
gathering, etc. (Massengill and Shuford, 1964).

13. Task analysis

Task analysis is discussed in Shuford's paper on degradation and optimal
learning (1963), toward the end of Toda's initial contract paper on the
F-E (1963a), and in Massengill and Shuford's paper concerning analysis of

performance (1964). For specific applications of task analysis, see
section 3 of this chapter, on the formulation of tasks.

V. ABSTRACTS OF COMPLETED REPORTS

Toda, Masanao. The optimal strategies in some simple fungus-eater games.
ESD-TDR-63-406. iii + 87 Pp., including 8 figures, 1 table, references.
July 1963a.

This paper discusses, mathematically, which strategy, where strategy
is a set of rules for making decisions, is good and which is bad under
what conditions in some one-person games. One purpose of the paper is
to uncover some of the basic logic needed to handle decision processes. A
second purpose is to provide an experimental context which will facilitate
the subject's understanding of the nature of sequential decision tasks.

The basic structures of discrete fungus-eater, F-E, games are described
in a section dealing with the definition of concepts and the classification
of games. Psychological implications of F-E games are mentioned. Of special
importance is the exposition of the feature in the F-E approach, the means
object, which is not dealt with in traditional psychological research. The
means object in F-E games is fungus, which represents objects or outcomes
which are primarily means to an end. Thus F-E games involve means objects
as well as end objects.

A section is devoted to optimal strategies in F-E games. The optimal
strategy of a given game is defined as the strategy which maximizes the
expected Uranium return. It is called the decision function and is a
function of the permanent, the external, and the internal decision contexts
of the game. Four theorems, illustrating the notion of interplay between
the characteristics of the environment and the F-E in the determination
of the optimal strategy, are proved. A standard procedure for solving
optimal decision functions is explained.

Optimal strategies are then derived for two infinite F-E games, G1 and G2, and one finite game, G3. And though the complete optimal solution is not given in this paper, some of the obvious characteristics of the solution are given for a second finite game, G4. The complete solution is given in (Nakahara and Toda, 1964a).

The results of a pilot study involving the first three games are reported. These findings show that few subjects were optimal except in the finite game. But they reveal a clue as to the nature of preferred human strategies and this furnishes encouragement to find the constraints that make this type of strategy optimal.

The author comments that the element, fungus, creates a new field in experimental psychology. It is a reward but its worth is purely conditional. He also comments on the importance of knowing optimal strategies before conducting experiments; namely that if the optimal strategy is not known, there is no satisfactory way of analyzing subjects' behavior. It is also desirable, in order to more fully understand subjects' performances, to know the strategies which are optimal under various constraints.

Finally, the author points out that the F-E approach is a research strategy. He compares this approach with the most popular approach in contemporary experimental psychology.

Toda, Masanao. Measurement of subjective probability distributions. ESD-TDR-63-407. iii + 42 Pp., including 13 figures, references. July 1963b.

Five experimental one-person games are derived which promise both valid and efficient measurement of subjective probability distributions, where subjective probability represents an individual's expectation as to the occurrence of an unknown event. The author gives four requirements that a desirable measurement procedure for subjective probability should satisfy. 1. The experimenter should understand the logical nature of the task presented to the subject. 2. The task should involve well-defined payoffs. 3. The structure of the task should be such that it is to the disadvantage of the subject to respond in a manner inconsistent with his expectations. 4. The technique should not be inconsistent with decision theory. The games described in this paper satisfy these requirements.

In section one, three games are derived for the measurement of discrete subjective probability distributions. The author defines the matching property which is used to obtain payoff rules and which leads to three classes of payoff functions used in measuring discrete subjective probabilities. Since each method may have its own psychological bias, a theorem is derived which may be used to minimize the bias by blending any two independent functions which have the matching property. The three games applicable to discrete situations: the spherical gain game, the logarithmic loss game, and the quadratic loss game, are derived. The author then describes displays which can be used to realize these games in experimental situations. Pilot studies done on the second two games are mentioned briefly.

In section two, two methods are derived for measuring continuous subjective probability distributions. In the first, called the range betting method, the subject bets on the true state in the form of a range. Two payoff functions are defined. Each can be used to estimate the unknown parameters of the distribution under the assumption that the subject is maximizing expected payoff. The range betting method can be applied to any subjective probability distribution, bounded or unbounded. The problem of non-unimodal cases is discussed. The range betting method is extended to discrete distributions. A second method for measuring continuous distributions, the partition betting method, is derived and compared with the range betting method.

Finally, the author points out the importance, when using these methods experimentally, of making the nature of the optimal solutions easy to understand and making the experiments as exciting as possible.

Toda, Masanao. Microstructure of guess processes: part c. ESD-TDR-63-548. v + 44 Pp., including 18 figures, references. Sept. 1963c.

An experiment is described in which human subjects predict the occurrence of one of two events in a complex sequence of binary events. This is a type of probability learning situation using what is known as the two-arm bandit or guessing experiment. The non-optimal strategy of probability matching, which has been found in numerous studies, is discussed. Then, trial-to-trial changes in the proportion of human subjects predicting one of the two events are analyzed, using operator reinforcement models of Estes and Bush and Mosteller. The direction of change predicted by these models is wrong on about 75% of the trials.

The author then develops three models and uses them to analyze the data. The models are called: the no-learning model, the time-dependent decay model, and the cycle-dependent decay model. The author believes that these provide some insight into the nature of probability learning. One conclusion is that the apparent simplicity of averaged guessing curves is a complete deception.

The author next turns to the matter of estimating parameters of stochastic processes and examines the question of what the next best procedures are when it is not possible to use optimal procedures. These next best procedures are not given in textbooks so the author attempts to derive them. Several procedures are compared. The author describes and points out the absurdity of one widely used estimation method, the method of simple sum. He then attempts to obtain a set of criteria for the admissibility of suboptimal methods of parameter estimation. Two criteria

are derived from the nature of the error function of the least squares method.

Finally, the method of minimum absolute error is recommended as being very useful. It estimates parameters by minimizing the sum of absolute errors. It disregards exceptional data values. It does not innocently give estimated values no matter in what way exceptional values may exist in the data, but on the contrary, gives precise information, through the course of estimation, about which values are exceptional and in what way.

Toda, Masanao and Shuford, E. H. Utility, induced utilities and small worlds. ESD-TDR-63-662. iv + 40 Pp., including 3 figures, references. Oct. 1963.

Some empirical findings concerning decision theory seem to suggest violation of the principle of maximizing subjectively expected utility. The authors believe that these "inconsistencies" can be explained in terms of the present inadequacy of decision theory to provide sufficient criteria for the unique formulation of decision tasks. This inadequacy may lead to disparate interpretations of the same task by different people. The authors discuss the empirical findings and attempt to show how they are accounted for by the inadequacy of decision theory; thus holding intact the principle of maximizing subjectively expected utility.

The first of these findings concerns the measurement of utility. It is shown that the empirical approach to measuring utility in one situation and using the results to predict behavior in another situation is based on the supposed constancy of personal decision contexts. But the empirical approach gives no reason to suppose that the personal context will remain constant. And further, it does not provide a rationale for generalizing from one personal decision context to another. The authors suggest that emphasis be placed on understanding the personal decision context. This may provide insight into the matter of the constancy or lack of constancy in personal decision contexts.

The authors then use an example from the game of chess to show that typical choice experiments do not measure utility but induced utility. Induced utility is defined as conditional expected utility, i.e., the expected utility a subject anticipates by possessing some object or attaining

some outcome, where the object or outcome is a means to some end. It is noted that induced utility changes with the relation holding between the object and the goal. Induced utility behaves in many ways like utility. In using induced utility, there is the same problem of generalizing from one personal context to another. The authors believe that analysis of means-end relations will help lead to the solution of this problem.

The material concerning the first set of findings, showing that decision theory is too narrowly conceived and needs extending, is concluded with a discussion of the problem of isolating small worlds.

The second set of findings is discussed under the heading of the "Chipman-Ellsberg-Fellner Paradox." This involves the seeming inconsistencies that appear in the empirical results of running subjects in a gamble involving two urns, one with known contents and the other with unknown contents. Separate interpretations of these results by Ellsberg and Fellner result in the suggestion that a third factor, in addition to subjective probability and utility, is needed to account for the observed choice patterns in the empirical results. The third factors suggested are related to the quality or quantity of information about the probabilities of the states. The authors show that a reformulation of the decision task in terms of a larger context will account for the empirical results without sacrificing the principles of maximizing subjectively expected utility.

The authors conclude that the major problems in the application of decision theory reside not in inadequacy in the principle of maximizing subjectively expected utility but in the difficulty of properly formulating decision problems. They suggest two ways of alleviating the difficulty. One is to improve techniques for temporarily closing open systems, i.e., extracting suitable small worlds from a grand world. The other is to

improve our ability to perceive how another individual formulates his

decision problems.

Shuford, E. H. Some Bayesian learning processes. ESD-TDR-63-623.
iii + 39 Pp., including 1 table, references. Oct. 1963.

Purposive mathematics, of which decision theory, information theory,
and dynamic programming are examples, can be used to provide a logical
framework for the analysis of tasks and can specify the best performance in
a task given the structure of the situation and the purpose of the user.
This paper introduces the technique of degradation in a decision theoretic
context as a method of logical analysis of tasks. The results given by
the method provide insight into the nature of certain learning tasks.

A well-defined two-act, two-state decision task, i.e., one with known
prior and conditional probabilities and utilities, in which data is observed
on each trial and the actual state obtaining is given after each trial, is
degraded in four different ways. While no learning is exhibited with
repetition of the well-defined task, the optimal strategies for the degraded
tasks exhibit optimal learning.

In the well-defined task, only the knowledge of the actual state
obtaining on a given trial is unknown to the decision maker at the time the
deicision is made. The first degradation is obtained by withholding the
values of the prior probabilities. But over trials, the prior probabilities
may be learned by using the feedback of the actual state obtaining on each
trial. Thus, available information is used to reduce uncertainty about the
prior probabilities. This is an example of learning prior probabilities
when payoff is given on each trial.

The second degradation is obtained by withholding the post-decisional
information as to the actual state as well as the prior probabilities. In

this case, the data observed prior to the trial on which the decision is being made is used in the optimal learning strategy to reduce uncertainty about the prior probabilities. This is a case of learning prior probabilities when payoff is delayed.

The third degradation is obtained by eliminating the knowledge of the conditional probabilities. The strategy for optimal learning in this situation can be characterized as learning posterior probabilities.

The fourth degradation is obtained by elminiation of all data. The decision maker must, in effect, predict the occurrence of the states. After each prediction, he is told which state actually occurred. The optimal learning strategy for this situation may be characterized as learning a simple probability.

For each of the degradations, the author derives the optimal strategy and gives several applications. Two interpretations of the set of degradations are offered. One is that each optimal strategy corresponds to the performance of an ideal decision maker operating under certain constraints. The author, at this point, makes a distinction between strict optimality and constrained optimality. The other interpretation is that each degradation represents an experiment which could be performed in the laboratory and for which the corresponding strategy is strictly optimal. This interpretation leads to a partial ordering on a set of possible decision tasks, the order being from the most structured to the least structured. Comparison of the expected payoff of the strategies for two tasks can yield a measure of the structure in the degraded task.

Finally, the author points out that some of the learning strategies mentioned above contain learning processes which may occur at two different levels. He calls the lower level process a learning process and the higher level one a learning-to-learn process.

Toda, Masanao and Shuford, E. H. Logic of systems: introduction to the formal theory of structure. ESD-TDR-64-193. v + 78 Pp., including 7 figures, 2 appendices. Jan. 1964.

This is a first in a series of papers. The purpose of the series is to illustrate the importance and necessity of using a research strategy based on the logical analysis of structure in the science of behavior, where the science of behavior includes all research concerned with the prediction and control of purposive systems. The authors discuss the reasons such a strategy is necessary and contrast it with the present dominant strategy in the study of behavior. This is followed by the development of a formal theory of structure.

The traditional research strategy is to search for basic relations and then attempt to construct more complex relations from them. The authors show that the conditions for success of this strategy are not met in psychology. They further show that one does not need to know all of the conceivable elementary relations but just those necessary for that structure which is relevant to the experimenter's purpose. Thus, searching only for basic relations may lead to the discovery of many irrelevant relations. They believe that the search for basic relations should be continued, but that this search alone is not sufficient for the purpose of establishing a full-fledged science of behavior. The authors conclude the introduction by recommending a search for the logic of structures, even though little is known at present about structures of specific situations.

The authors then proceed to develop a formal theory of structure. This includes a discussion of two types of decomposition, locational and functional, as well as a discussion of composition. The state, as opposed

to the structure of a system, is discussed. Relations between systems are distinguished in terms of quantitative relations and categorical relations. Primary structure and hyperstructure are defined. A distinction is made between atomic state descriptions and non-atomic state descriptions. Open and closed systems and the independence between systems are also discussed.

An application of the theory to psychology is sketched in a final remark. It gives an example of possible conceptual processes.

Appendix I contains a glossary of terms used in the text, as well as references to the pages in the text on which the terms are defined. Appendix II contains the theorems, corollaries, and the axiom used in the text.

Nakahara, Jun-ichi and Toda, Masanao. Optimal strategies and human behavior in fungus-eater game 4. ESD-TDR-64-237. iv + 60 Pp., including 40 figures, references. Feb. 1964.

The optimal strategy for fungus-eater game 4, G4, is derived and a pilot experiment involving human beings performing in this game is reported. Though this paper assumes familiarity with the preceeding papers in this series, there is a brief summary of the points in the preceeding papers that are relevant to G4. In the summary, discrete F-E games and their optimal strategies are described. G4 is characterized as a binary, homogeneous game with a vision-span, V-span, of 1. The three variables specifying the state of the F-E: F-storage, U-storage, and L-storage, are defined and described. A definition of the well-informed F-E is given and the permanent decision context, the external decision context, and the internal decison context are defined. The optimal decision function is defined as the one which maximizes future U-return. There is a description of the IC (Internal Context) diagram, its absorption barriers and critical levels. A point on the IC diagram represents the internal context of the F-E. The four types of unit environment of G4 are given, together with an explanation of their probabilities.

In discussing G4, it is pointed out that only one of the four possible external contexts is relevant to the decision function. The explicit expression for the optimal decision function is written in terms of G4. A delta function is defined which specifies the value of the optimal decision function for a given internal decision context. The decision shifting point is described. The decision shifting point, for each line below the critical level, divides the IC diagram into two regions, the U

- 64 -

decision region and the F decision region. Thus, the optimal strategy for G4 can be found by obtaining the decision shifting point for each line below the critical level. The decision for the critical level, y=0, is to always take U. An analytic solution for y=1 is obtained and an outline for the solution for y=2 is given.

Since it is very difficult to get analytic solutions for values of y greater than 1, an approach which uses the computer to obtain numerical solutions is adopted. The recurrent relation of the expected U-gain function is given and the numerical solution for obtaining the optimal solution is described. A flow chart for the computations is included. The computer program is used to obtain the numerical solution, describing the border between the F and U decision regions, for various values of the parameters involved. The results are plotted on IC diagrams. Plots of the values of the maximum expected U-return for various internal contexts are also shown.

The last section includes graphs showing the locomotion line of subjects who participated in a pilot experiment involving G4. Decisions and predecisions of the subjects are shown and compared with the border describing the optimal decision. The experimental procedure is described. While the data is insufficient to uncover the decision strategies actually employed by the subjects, the authors conclude: 1. that most of the subjects used a critical x-value strategy, which varied from subject to subject and with parameter values, and 2. that the observed strategies were not far from optimal.

VI. ABSTRACTS OF REPORTS IN PREPARATION

Massengill, H. E. Pre- and post-experimental analysis of truncated sequential sampling situations. In preparation. 1964b.

This paper contains the symbolic programs and flow charts of computer programs for the pre- and post-experimental analysis of truncated sequential sampling situations as well as an explanation of the theory used in designing them and step-by-step instructions on how to use them. The programs described are designed specifically for three-act, where one act is "to sample," two-state decision situations with cost of sampling linear in n, the number of sample returns observed. But the basic structure of the program is designed so that analysis is possible for situations with any number of acts and states and any payoff function; subject only to the condition that the sampling process be binomial. Alternate analyses can be performed by changing relevant subroutines in the main program.

Both the pre- and post-experimental analyses are written in the Decal symbolic language for the PDP-1 computer. The pre-experimental analysis can be run in one stage, using a minimum of two cores, or in two stages, using one core. Given this amount of storage, situations with a truncation level of up to 50 can be handled. The post-experimental analysis uses only a small section of one core.

Phase 1 of the pre-experimental analysis works down the decision tree, where the top level is the truncated level. At a given level of the tree, the expected losses of each act for each n, r combination of that level, where r is the number of successes obtained in n observations, are obtained. The losses at that level may be punched out for future reference after the level has been analyzed. Whether or not the losses are punched out, they

are destroyed, in order to save storage, as the losses for the next level are computed. For a given n and r, if it is optimal to sample, a "1" replaces a "0" in a given bit of storage location. The "0" remains in the bit if it is not optimal to sample. This information is sufficient to make phase 2 of the pre-experimental analysis possible.

Phase 2 uses the 1's and 0's from phase 1 to go back up the tree in order to find each optimal path. The output of phase 2 is a frequency distribution of the possible optimal paths, a probability distribution of the paths, the expected length of the optimal paths for that tree, and the expected loss incurred by sampling at the lowest level of the tree (which is obtained in phase 1). These latter two numbers characterize the tree in question.

One purpose of the pre-experimental analysis is to provide specific knowledge concerning the nature of sequential sampling. A second purpose is to aid in choosing parameters for experiments in which the underlying task is truncated sequential sampling. A third purpose is to provide information concerning optimal performance for use in the analysis of performance. This provides a standard against which to compare the behavior of human beings performing in the task in order to determine the least upper bounds of their capacities in this task.

The program for post-experimental analysis simply computes the expected loss of a subject in a truncated sequential sampling task after each trial, keeps a running sum of these losses during the experiment, and averages over the trials at the end of the experiment. This average is taken as an estimate of what the subject would do if he were to perform in the situation, characterized by this tree, many times. The program outputs this number, along with the expected performance of the optimal strategy

for this tree. The latter number is obtained during the pre-experimental analysis and is the minimum expected loss of the three acts before any observations have been taken. Analogously, the sample size taken by the subject on each trial is averaged over trials and compared with the expected optimal sample size for the tree.

Massengill, H. E. A study of conditional probabilities. In preparation. 1964b.

Decision theory, from the point of view of the everyday decison maker, can be viewed as a tool. The majority of a person'a decisions can be made intuitively and the outcomes of these decisions are largely what would be expected. It is not worthwhile for the decision maker to go into a deeper analysis of these situations because to do so might cost more than the analysis is worth in terms of net gain in expected value. However, there is a class of decisions that a person must make that, while small in number, is of sufficient importance to justify paying the cost of thinking involved in making a deeper analysis. It is in these situations that decision theory can be of use to the decision maker in problems of everyday life.

There are two uses to which decision theory can be put in these situations. One is simply to help the decision maker organize his thoughts on a problem and get a more specific formulation of the matters involved. Another is to give exact solutions in situations in which the decision maker wants to specify his probabilities and utilities but doesn't want to combine them himself for a final decision or wants to check his own decision against the exact solution. This paper concerns the former aspect, i.e., that of using decision theory as a tool to help clarify decision situations which the decision maker feels are worth clarifying.

Two ways in which a decision maker might be able to improve his formulations of decision situations are discussed. The first concerns the classification of information in decision problems in terms of the categories specified by decision theory, e.g., states of nature, outcomes, etc. The second concerns the matter of deciding what kind and how much, if any,

further information should be sought. Both involve the use of conditional probabilities, our main interest in this paper.

There are some decision situations in which the process that is generating information is known to such an extent that the conditional probabilities can be derived logically. But in many everyday situations, this is not the case. This brings up the matter of using an empirical approach. An extended example is given of the formulation of a problem involving empirical conditional probabilities. This example concerns a theatrical producer who wishes to use decision theory as an aid in the formulation of his problem of whether or not to produce a certain play. The problem is dealt with in terms of classifying what is already known into the categories specified by decision theory as well as deciding whether or not to get more information.

Next, there is a consideration of the processes which generate information. First, there is a discussion of the effect on performance of not fully understanding the process which is generating the information. Reference is made to (Shuford, 1963) to illustrate that even when the data-generating process is not fully understood, there may be enough information present so that, asymptotically, performance can approach that possible when the information generating process is fully understood. But the difference in performance, in the early trials of some situations, points up the significance of fully understanding the process in these situations. Second, the information generating processes are classified according to whether they are deterministic or probabilistic, with the latter being subdivided into stable and unstable processes. The significance of the information generated by these various types of processes is examined.

Finally, there is a discussion of the relation between conditional
and prior probabilities. When conditional and prior probabilities are
independent, given conditional probabilities can be used with various prior
probabilities in order to obtain posterior probabilities. For example, the
producer of a play may have to make the decision to produce or not produce
a play for many plays. And though different plays are involved, some of
the information relevant in one situation may also be relevant in others. If
this information were imbedded in posterior probabilities, instead of being
in the form of conditional probabilities, it would be more difficult to use.
Since there are occasions when it would be helpful to have conditional
probabilities that are contained in posterior probabilities, procedures
for extracting them are introduced.

Massengill, H. E. Teaching optimal behavior in a computer-controlled percentage estimation task. In preparation. 1964c.

This paper describes the attempt to use the computer to teach optimal behavior in a well-defined decision task involving percentage estimation. The relation of programmed instruction to the design of the experiment is discussed. Following an explanation of the experimental procedure, there is a detailed description of both a trial and a session, showing what is the same and what is different in each for each of the three groups of subjects involved.

The experiment is analyzed from two points of view. One concerns the task involved in the experiment and the other the empirical data observed in the experiment. The analysis of the task is a decision theoretic analysis in which the task is examined in terms of various levels of optimal performance possible in the task, given certain assumptions concerning the information available in the task. The value of the various types of information in the task is discussed. The actual information that the subjects had is also discussed, e.g., they were given the outcomes of the data-generating process but not the probabilities of the outcomes.

The data analyses are of two types, one traditional and the other decision theoretic. The traditional analysis is an application of t tests between the various groups. The decision theoretic analysis consists of two types of comparisons. One is a comparison of the actual earnings of groups and of individuals with the earnings of the optimal strategy for various possible levels of optimal performance. The other is a comparison of the responses of individual subjects with the optimal responses. This

includes a regression analysis and an estimate of each subject's prior parameters, $\underline{r}'$, $\underline{n}'$, for each decision context, given that they used the sample information available, i.e., $\underline{n}$, the number of observations, and $\underline{r}$, the number of successes. There also is an estimate of each subject's $\underline{n}$, for each decision context, given that he used the actual $\underline{r}'$, $\underline{n}'$, and $\underline{r}$.

Suggestions are made concerning various modifications of the experiment which might yield additional information concerning the performance of subjects in estimation tasks. A decision theoretic analysis for estimation experiments, in terms of expected earnings rather than actual earnings, is outlined.

Two applications of the experimental work described in the paper are discussed. One is an application to situations in which an employer is seeking to determine which prospective employees to hire for a decision making task, where the choice is between men, men and machines, and man-machine combinations. The other is a discussion of the results of the experiment in terms of implications for a suggested approach to teaching intuitive decision making.

Appendix I is a transcript of the instructions given a typical subject including the variations applicable to the different groups involved. Appendix II is a technical note concerning the discovery of what is believed to be the analytic solution for the median of beta distributions. This is important for the situation described in this paper because the prior and posterior distributions relevant to it are beta distributions and the optimal response for any trial is the median of some beta distribution. This discovery of the analytic solution of the median of beta distributions makes possible an analytic solution for the expected value of the optimal strategy, for a given decision context, in this experiment. Appendix III

gives a listing of and brief description of the various computer programs developed for the decision theoretic data analysis of this experiment.

Massengill, H. E. and Shuford, E. H. The analysis of performance. In preparation. 1964.

Analysis of performance is especially important in those applications of decision theory in which a person is asked to adopt someone else's formulation of a task. For example, an employer may desire to select employees to perform a decision task. In such cases, it is important to have some way of determining how well prospective employees are able to do in a given task so that information will be available on which to base a decision concerning which, if any, to hire. The employer may also wish to consider training employees to perform tasks, given the company's formulations. In this situation, he needs some idea of how well the prospective employees perform without training and a way of evaluating their performance as they are trained. The employer may also be concerned with the matter of allocation of task components between men and machines and thus need to be able to compare the performance of the two.

In this paper, two cases are discussed in which an employer is considering using human beings, machines, or some combination of the two to make decisions for the company. The first case concerns a well-formulated task in which the employer knows the optimal strategy, though the strategy may change because of changes in the values of the parameters involved. He plans to compare human and machine performance for representative values of the parameters. In the second case, the task is very complex and not well-formulated. Thus there is no analytic solution or even numerical solution available for determining the optimal strategy.

The paper discusses what the employer can do in terms of evaluating the performance of human beings in both types of tasks as well as what

measures of performance to use. For the well-formulated task, the employer can test prospective employees and compare the performance of each with the strictly optimal performance, which is the performance that would result if a machine were used. Then he can compare each man with the machine in terms of the time required to perform the task, in terms of the resetting that might be required as the problem varies, etc. The costs can then be balanced off against the differences in performance. The logic, of course, is to pick the best men and/or machines. There is no final answer as to which are best, but we can get some idea of what can be expected. The well-formulated situation is illustrated with examples from an estimation task and from a sequential sampling task. The value of information is also discussed. For example, certain information in a given situation may contribute so little to the optimal solution, in terms of the cost of using it, that it can be ignored.

In the second case, in which the task is not well-defined, enough is known so that the nature of the task can be specified, e.g., it may be dynamic programming, sequential information gathering, etc. A well-formulated task of that class can be used to test the man. This, of course, is what is done in all testing. And the man's performance can be compared with machine performance in the well-formulated task. In order to gain insight into the matter of performance in the version of the task that is not well-formulated, suboptimal procedures are used and the resulting performance studied.

In this type of application, the measure of the subject's performance is taken as an indication of the least upper bounds of his capacity in the task. The measure of performance that is recommended is the expected value of each act chosen, rather than the actual earnings for each act. The use of this measure reduces variability.

Nakahara, Jun-ichi and Toda, Masanao. Simple discrete fungus-eater games of class III: the infinite co-existent V-span 2 chain game. In preparation. 1964b.

A brief summary of the structure of discrete F-E games and their optimal strategies is given. The reader is assumed to be familiar with the two preceeding reports in this series. The game discussed in this report, the infinite co-existent V-span 2 chain game, G1-2, is identical with G1-1 except that in G1-2, the F-E's V-span is 2 rather than 1 and the branch structure of the environment is explicitly assumed to be a chain.

Because the V-span is 2 and the chain structure is co-existent, there are four kinds of external decision contexts: [FU] - [FU], [FU] - [OO], [OO] - [FU], and [OO] - [OO]. The permanent decision context is described by the probability that FU will occur at a choice point, p, and by the parameter which represents the increase in F-storage when one F is taken, a. The F-E's F-storage is the only relevant parameter in the internal decision context. For G1-1, the optimal decision function for the three cases: pa<1, pa=1, pa>1, were discussed. This report is restricted to the optimal decision function for the case, pa≤1, because it can be easily shown that when pa>1, future U-return can be made infinite.

Section three introduces the theorem of mortality which states that if the branch structure of the F-E's world is chain and if pa=1, in an infinite co-existent game, then the F-E is mortal. The authors prove a lemma which states that the strategy which gives the longest life to F-E, in an infinite co-existent chain game, is strategy F, always take fungus. Thus to prove the theorem of mortality, it is only necessary to show that the F-E is mortal even under strategy F. The gambler's ruin problem is used to aid in this proof.

For pa<1, the optimal strategy in Gl-1 is strategy U, always take uranium. But this is not the case for Gl-2 since, for example, there is a situation in which F should be taken, followed by taking U. The strategy, $\tilde{U}$, which says to take F when the F-storage is one and the external context is [FU] - [FU] and to take U otherwise, is discussed in detail. It is concluded that the strategy $\tilde{U}$ is optimal if pa<1.

Nakahara, Jun-ichi, Shuford, E. H., and Toda, Masanao. The learning
fungus-eater in G3. In preparation. 1964a.

This paper discusses the adaptive version of the fungus-eater in
game 3, G3. The well-informed version of G3 is described in a previous
report in this series. The task of the F-E in both versions is to seek
the optimal policy under uncertainty. The learning F-E in G3 differs
from the well-informed F-E in G3 in that he is uncertain about the proba-
bility of the occurrence of the external decision context E(F U), $\underline{p}$, on a
unit environment.

Following the introduction, the extensive analysis of adaptive G3
is discussed. The observation process of G3 is shown to be a Bernoulli
process. The F-E is assumed to have a prior probability density function
which, taken in conjunction with the observation process, results in a
posterior distribution. Since the variance of the posterior distribution
can be expected to be smaller than that of the prior, the uncertainty of
the parameter, $\underline{p}$, can be reduced progressively during U-searching trips.

The F-E's internal decision context can be represented by four
parameters. When an internal context is followed by a chance move,
external context E(F U) occurs with probability, $\underline{p}$, and context E(O O)
occurs with probability, 1 - p. Thus, there will be three actions open
to F-E: take nothing when E(O O) occurs, take F when E(F U) occurs, or
take U when E(F U) occurs. The game tree representation of a choice point
is shown.

The F-E game tree, the tree representing all of the possible courses
of action by the F-E is then discussed. Expected future U-return can be
assigned to the nodes of the tree. Thus the optimal strategy is available

for selecting the action, at each choice point, which leads to the greater expected future U-return. The process of determining the optimal policy consists of generating the whole game tree and the expected U-returns and then tracing backwards in terms of expected U-return.

Section three discusses the computational method used to deal with the adaptive version of G3. Two operators, the game tree generating operator and the expected future U-return induction operator are introduced and explained. Their purpose is to formalize and simplify the F-E game tree generating process and the expected future U-return induction process, respectively. This, in turn, helps to simplify the conceptualization of the computer programming necessary for determining the optimal policy. A flow chart showing the programming steps for a sample problem is given.

It is pointed out that when the F-E's life span left for future travel, given by the positive integer N, is large, the computational problem becomes very difficult, since the possible number of different courses of action can be 3^N. To obtain solutions for large N's on computers with limited capacity, in a limited amount of time, it will be necessary to develop methods of approximation. Such methods will be discussed in future reports.

Nakahara, Jun-ichi, Shuford, E. H., and Toda, Masanao. The accident-prone fungus-eater in G3. In preparation. 1964b.

This paper introduces another process, the accidental stopping or accidental death process, into discrete F-E games in which the player is well-informed. In this paper, the process is applied to finite game G3. Familiarity is assumed with previous work on the F-E and appropriate references are given.

In the accidental stopping process, the F-E may be forced to stop the game on each trial with some fixed probability, π, independent of the probabilities of the external contexts, though he has not exhausted his F-storage or L-storage. Thus, a probabilistic absorption barrier is added to the two deterministic barriers, the dooms-day barrier and the starvation barrier, previously introduced. The F-E is assumed to be well-informed in these games, i.e., he knows the probability of accidental death as well as the values of the other parameters.

Two variations of the accidental death process are introduced. In one, called the pre-substance case, the accident, if it happens, happens before the F-E can collect what he has chosen at that choice point, e.g., before he can collect uranium if he has chosen uranium. In the other, called the post-substance case, the accident, if it happens, happens after the F-E has collected what he has chosen.

For both the pre-substance and the post-substance cases of accident-prone G3, the authors derive the decision function, the resulting decision rule, an explicit expression for the expected U-return function, and the boundary equations which make numerical solution by simple iteration possible. Examples of the pre-substance and post-substance cases for

for various values of the relevant parameters are included. This is followed by a general discussion of the accidental death process and a comparison of the optimal solution in accident-prone G3 with the optimal solutions in G1, G3, and G4.

VII. SUMMARY AND RECOMMENDATION

It should be clear by now that the two central concepts of our approach are purpose and optimality. We assume that systems composed of individuals performing tasks are purposive systems. And we assume that human beings who perform in these tasks behave optimally given their formulations and constraints. The former assumption implies that we need a language of purpose in order to be able to talk about the systems. This, along with our desire to begin with a theory which is comprehensive, which is internally consistent, and which will allow us to map behavior onto logic and mathematics, has led us to choose decision theory as the basis for our approach. The second assumption that we make, the rationality assumption, allows us to accomplish the desired mapping. These considerations put us in a position to proceed with the development of a science of behavior.

Since our goal is to explain all of the important aspects of behavior, we feel that it is necessary to begin with a comprehensive theory rather than with miniature theories and to study only the significant tasks of human beings. Since we can begin with a comprehensive theory which is internally consistent and which will remain so throughout its development, we feel that we have an advantage over those who begin with miniature theories with the hope that one day they will all fit together into one grand comprehensive theory. We feel that we have the advantage because we are reasonably sure that the fitting together of miniature theories will not yield an internally consistent theory because the rules of structure for fitting them together are not known and/or because

they are based on the arbitrariness and instability of human behavior.

Since decision theory is a branch of mathematics, we can develop it by logical deduction without resourse to empirical observation. Our procedure, which we call task analysis, is to decompose and then recompose the significant tasks that people perform in order to find the optimal behavior involved. In this way, we can build up a catalogue of formulation-constraint combinations with the resulting optimal behavior associated with each. Ultimately, we will be able to look at the behavior of an individual in various tasks and infer his formulations and constraints, i.e., understand his behavior.

Because there has been considerable misunderstanding concerning the descriptive and normative aspects of decision theory, we will summarize our feelings on this matter. It should be clear to those who have followed the arguments in this paper that, at this point in its development, decision theory cannot describe behavior in the sense of specifying formulations and constraints on the basis of observed behavior. Thus, it is meaningless to try to test decision theory as a descriptive theory. This does not mean that decision theory will never describe behavior in the above sense. It does mean that we will not be in a position to know whether it can or not until it has been developed to the point where tests are meaningful. And we must not delude ourselves about the extent of development that will be necessary for such tests nor about the time that this development will take. The extent is great and the time will be long.

Concerning the normative aspects of decision theory, we have seen that we cannot formulate a task, find the optimal behavior for a given purpose, and say that this represents the behavior that should be given by everyone who wants to behave optimally in the task. It is meaningless

to talk about how a person should behave in a task, if his formulation, purpose, and constraints are not taken into account.

In applications, decision theory can be used to set up standards of behavior in situations in which a person is asked to adopt another person's formulations. This does not mean that the standard represents the optimal behavior for any person performing the task but only for those who adopt the formulation and purpose involved and who have the constraints that are implied. But we can use the standard to comment on the quality of a person's performance, when the person has sought to adopt the formulation and purpose specified, and we can get some idea as to whether his performance in the task can be improved.

In our discussion of decision theory, we should not forget its function as a tool. Decision theory, as an area of applied mathematics, is useful as a tool in analyzing tasks. This analysis may be performed with the view of developing a science of behavior or with the view of improving performance. The latter aspect is the key to our approach in applications. We should make it clear that the decision theory can be accepted as a tool without accepting it as ultimately being a descriptive theory of behavior. This is because decision theory is a normative theory in the sense that for a given formulation, purpose, and set of constraints, it does specify optimal behavior. And this is all that is necessary for applications to be possible.

Our recommendations are implicit in the body of the paper. In basic research, we recommend that the development of a science of behavior be approached by the analysis of tasks rather than by the empirical study of human behavior and/or the attempts to set up miniature theories. This approach involves the assumption that systems composed

- 85 -

of human beings performing tasks be accepted as purposive systems and the assumption that people maximize subjectively expected utility, i.e., behave optimally, given their formulations and constraints. If one is willing to accept these assumptions and is willing to accept decision theory as the basis for the approach, then he is ready to proceed in the attempt to develop a science of behavior by the analysis of tasks.

In applied research, we recommend that the results of task analysis be used to help improve intuitive decision making, to help decision makers get exact solutions for their more important decision problems, and to aid in the development of decision systems so that overall performance can be improved. As we have said, one can act on these recommendations in applied research without accepting our rationality assumption. It is sufficient to assume that people can learn to behave optimally in some tasks or, at least, that the results of task analysis are meaningful in designing machines to perform certain important tasks.

The everyday decision maker can look upon our efforts as an attempt to develop tools to help him improve his performance in decision problems. We assume that performance in decision situations can be improved, not because people do not now behave optimally, but because there are certain constraints involving information and capacity that might be relaxed somewhat by training the decision maker or by helping him to get exact solutions. The greater the development of decision theory, the lower will be the cost of utilizing it both in terms of improving intuitive decision making ability and in terms of helping decision makers get exact solutions (Shuford and Organist, 1964).

As a final recommendation, we would like to suggest that the implications of our overall approach be examined in terms of possible

applications. This has been done, for instance, in S-R psychology. The idea of using small steps in programmed instruction probably developed not only out of the S-R theory, in which the association of a simple stimulus and a simple response and the resulting reward for making such an association are important, but also out of the general approach of dealing with elementary relations and feeling that once these elementary relations are known, the more complex relations will emerge.

It may be that an examination of our approach, in terms of the importance of choice and in terms of our emphasis on the decomposition and composition of significant tasks, may lead to a new concept of what should be taught and how it should be taught. For example, it may be that letting a person work with more complex systems and letting him decompose and compose such systems, might be more effective than decomposing systems for him with the possible result that he fails to see the structures of the systems or that he supplies irrelevant structures. This, of course, is merely a suggestion of a possible application of our overall approach and must stand or fall according to its effectiveness. But it is an example of one of the ways in which our approach might be applied to significant problems that are not handled directly by our theory.

Finally, we should include a word of caution to any who might be tempted to adopt our approach. It is true that decision theory has been developed to the point that it can be used to handle significant problems from the standpoint of applications. In fact, one of our most important tasks from this standpoint is to train people to use decision theory and to make it easier for them to use. But though the picture is bright in applications, it is not so bright from the standpoint of basic

research. Though much work has been done on decision theory, the techniques available at this time are still very limited. We are not yet in a position to handle the significant tasks that we believe must be studied if a science of behavior is to be developed. For instance, most of the work concerning constraints operating in purposive systems has had to do with informational constraints. But we must begin to consider very seriously constraints on capacity, e.g., those having to do with computing ability and memory, if we are to make important advances in basic research.

But though we have only begun in our attempts to develop decision theory, we have seen a steady improvement in our ability to handle more complex problems. As Shuford comments in his paper on Bayesian learning processes (1963, p. 1), the tasks that a few years ago were considered as much too complex to permit a logical analysis are now perceived as relatively simple and we use the term "complex" to describe more difficult tasks. We expect this process to continue until we hesitate to call any task complex.

REFERENCES

Massengill, H. E. Pre- and post-experimental analysis of truncated
   sequential sampling situations. In preparation. 1964a.

Massengill, H. E. A study of conditional probabilities. In preparation.
   1964b.

Massengill, H. E. Teaching optimal behavior in a computer-controlled
   estimation task. In preparation. 1964c.

Massengill, H. E. and Shuford, E. H. Analysis of performance. In preparation.
   1964.

Nakahara, Jun-ichi, Shuford, E. H., and Toda, Masanao. The learning
   fungus-eater in G3. In preparation. 1964a.

Nakahara, Jun-ichi, Shuford, E. H., and Toda, Masanao. The accident-prone
   fungus-eater in G3. In preparation. 1964b.

Nakahara, Jun-ichi and Toda, Masanao. Optimal strategies and human behavior
   in fungus-eater game 4. ESD-TDR-64-237. Feb. 1964a.

Nakahara, Jun-ichi and Toda, Masanao. Simple discrete fungus-eater games
   of class III: the infinite co-existent V-span 2 chain game. In
   preparation. 1964b.

Shuford, E. H. Some Bayesian learning processes. ESD-TDR-63-623. Oct. 1963.

Shuford, E. H. and Massengill, H. E. Some Bayesian learning processes:
   normal data generating process. In preparation. 1964.

Shuford, E. H. and Organist, W. E. Lowering the cost of analytic rationality.
   Paper read at the Conference on Bayesian Systems, U. of Mich. May 1964.

Toda, Masanao. The optimal strategies in some simple fungus-eater games.
   ESD-TDR-63-406. July 1963a.

Toda, Masanao. Measurement of subjective probability distributions.
   ESD-TDR-63-407. July 1963b.

Toda, Masanao. Microstructure of guess processes: part c. ESD-TDR-63-548.
   Sept. 1963c.

Toda, Masanao and Shuford, E. H. Utility, induced utilities, and small
   worlds. ESD-TDR-63-622. Oct. 1963.

Toda, Masanao and Shuford, E. H. Logic of systems: introduction to the
   formal theory of structure. ESD-TDR-64-193. Jan. 1964a.

Toda, Masanao and Shuford, E. H. Logic of systems: purpose, decision and
    structure. In preparation. 1964b.

DOCUMENT CONTROL DATA - R&D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author)	2a. REPORT SECURITY CLASSIFICATION
Institute for Research, State College, Pennsylvania	Unclassified
	2b. GROUP n/a

3. REPORT TITLE

Purposive Systems Theory and Application

4. DESCRIPTIVE NOTES *(Type of report and inclusive dates)*

Final Report

5. AUTHOR(S) *(Last name, first name, initial)*

Massengill, Edward H., Jr.

6. REPORT DATE	7a. TOTAL NO. OF PAGES	7b. NO. OF REFS
Jul 64	96	17

8a. CONTRACT OR GRANT NO.	9a. ORIGINATOR'S REPORT NUMBER(S)
AF19(628)2968	
b. PROJECT NO.	ESD-TDR-64-531
c.	9b. OTHER REPORT NO(S) *(Any other numbers that may be assigned this report)*
d.	

10. AVAILABILITY/LIMITATION NOTICES

Qualified requesters may obtain from DDC.
Aval from OTS.

11. SUPPLEMENTARY NOTES	12. SPONSORING MILITARY ACTIVITY
	Decision Sciences Labs., ESD, L G Hanscom Field, Bedford, Mass.

13. ABSTRACT

The purpose of this paper is to summarize the approach and theory on which the research performed under ESD contract AF19(628)2968 is based. Basically, the approach is the use of decision theory, with the assumption that people behave optimally given their formulations and constraints, to study the significant tasks that people perform. The ultimate goal of the approach is to map human behavior onto logic and mathematics. The emergence of the approach is given along with four basic requirements that we make of any theory to be used in understanding the behavior of individuals. Our approach is contrasted with more traditional approaches. The procedure of our approach, task analysis, is explained and is illustrated by examples from the contract research. The place of applications in our approach is dealt with extensively. The paper includes a guide to the more important ideas dealt with in the contract research with references to the relevant contract publications. Abstracts of these publications, seven completed and seven in preparation, are also included.

DD FORM 1473
1 JAN 64

14. KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Decision Sciences Theory of Probabilities Bayesian Systems Guess Processes Logic of Systems Purposive Systems						

INSTRUCTIONS

1. ORIGINATING ACTIVITY: Enter the name and address of the contractor, subcontractor, grantee, Department of Defense activity or other organization (corporate author) issuing the report.

2a. REPORT SECURITY CLASSIFICATION: Enter the overall security classification of the report. Indicate whether "Restricted Data" is included. Marking is to be in accordance with appropriate security regulations.

2b. GROUP: Automatic downgrading is specified in DoD Directive 5200.10 and Armed Forces Industrial Manual. Enter the group number. Also, when applicable, show that optional markings have been used for Group 3 and Group 4 as authorized.

3. REPORT TITLE: Enter the complete report title in all capital letters. Titles in all cases should be unclassified. If a meaningful title cannot be selected without classification, show title classification in all capitals in parenthesis immediately following the title.

4. DESCRIPTIVE NOTES: If appropriate, enter the type of report, e.g., interim, progress, summary, annual, or final. Give the inclusive dates when a specific reporting period is covered.

5. AUTHOR(S): Enter the name(s) of author(s) as shown on or in the report. Enter last name, first name, middle initial. If military, show rank and branch of service. The name of the principal author is an absolute minimum requirement.

6. REPORT DATE: Enter the date of the report as day, month, year; or month, year. If more than one date appears on the report, use date of publication.

7a. TOTAL NUMBER OF PAGES: The total page count should follow normal pagination procedures, i.e., enter the number of pages containing information.

7b. NUMBER OF REFERENCES: Enter the total number of references cited in the report.

8a. CONTRACT OR GRANT NUMBER: If appropriate, enter the applicable number of the contract or grant under which the report was written.

8b, 8c, & 8d. PROJECT NUMBER: Enter the appropriate military department identification, such as project number, subproject number, system numbers, task number, etc.

9a. ORIGINATOR'S REPORT NUMBER(S): Enter the official report number by which the document will be identified and controlled by the originating activity. This number must be unique to this report.

9b. OTHER REPORT NUMBER(S): If the report has been assigned any other report numbers (either by the originator or by the sponsor), also enter this number(s).

10. AVAILABILITY/LIMITATION NOTICES: Enter any limitations on further dissemination of the report, other than those imposed by security classification, using standard statements such as:

(1) "Qualified requesters may obtain copies of this report from DDC."

(2) "Foreign announcement and dissemination of this report by DDC is not authorized."

(3) "U. S. Government agencies may obtain copies of this report directly from DDC. Other qualified DDC users shall request through

_____ ."

(4) "U. S. military agencies may obtain copies of this report directly from DDC. Other qualified users shall request through

_____ ."

(5) "All distribution of this report is controlled. Qualified DDC users shall request through

_____ ."

If the report has been furnished to the Office of Technical Services, Department of Commerce, for sale to the public, indicate this fact and enter the price, if known.

11. SUPPLEMENTARY NOTES: Use for additional explanatory notes.

12. SPONSORING MILITARY ACTIVITY: Enter the name of the departmental project office or laboratory sponsoring (paying for) the research and development. Include address.

13. ABSTRACT: Enter an abstract giving a brief and factual summary of the document indicative of the report, even though it may also appear elsewhere in the body of the technical report. If additional space is required, a continuation sheet shall be attached.

It is highly desirable that the abstract of classified reports be unclassified. Each paragraph of the abstract shall end with an indication of the military security classification of the information in the paragraph, represented as (TS), (S), (C), or (U).

There is no limitation on the length of the abstract. However, the suggested length is from 150 to 225 words.

14. KEY WORDS: Key words are technically meaningful terms or short phrases that characterize a report and may be used as index entries for cataloging the report. Key words must be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location, may be used as key words but will be followed by an indication of technical context. The assignment of links, rules, and weights is optional